Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: August 18, 2012

G. Bertrand, Ed.
France Telecom - Orange
F. Le Faucheur
Cisco Systems
L. Peterson
Verivue, Inc.

February 15, 2012

Content Distribution Network Interconnection (CDNI) Experiments
draft-bertrand-cdni-experiments-02

Abstract

This document reports studies and related experiments on CDN interconnection performed by France Telecom-Orange Labs. The document summarizes implications of CDN interconnection to CDN service providers and lessons learned through CDNI experiments.

The main purpose of the experiments was to test the interconnection of CDN solutions from two vendors (namely, Cisco and Verivue) and to identify the gaps and needs for standardization work for CDN interconnection.

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 August 18, 2012.

Copyright Notice

Table of Contents

1. Introduction

This document reports studies and related experiments on CDN interconnection performed by France Telecom-Orange Labs. The document summarizes implications of CDN interconnection to CDN service providers and lessons learned through CDNI experiments.

The main purpose of the experiments was to test the interconnection of CDN solutions from two vendors (namely, Cisco and Verivue) and to identify the gaps and needs for standardization work for CDN interconnection.

This study is not intended to explore the entire scope of CDNI, and in fact, it purposely takes a minimalist approach. That is, we focus on what's minimally required to interconnect two cooperating CDNs in a "best effort" way. This provides a constructive foundation for adding requirements and mechanisms only after they prove essential in practice.

1.1. Terminology

We adopt the terminology described in [RFC3466], [RFC3568], [RFC3570], CDNI problem statement draft [I-D.ietf-cdni-problem-statement], CDNI framework draft [I-D.davie-cdni-framework] and CDNI use cases draft [I-D.ietf-cdni-use-cases].

Content Delivery Service

Set of services offered to content service providers (CSPs) for delivering their content through a single Content Delivery Network or interconnections of Content Delivery Networks.

2. CDN Interconnection Experiments

2.1. Experiment Configuration

The interconnection of two CDN solutions from different vendors has been tested. These tests have been run with CDN solutions from Cisco (hereafter referred to as Vendor A) and from Verivue/CoBlitz (hereafter referred to as Vendor B).

As depicted in Figure 1, we have interconnected two experimental CDNs (CDN A and CDN B) operated by different subsidiaries of a large CDSP. The CDNs lab equipment were located in two different countries, henceforth referred to as Country A and Country B and they relied on CDN solutions from two different equipment vendors, namely, Vendor A

and Vendor B. The CDNI experiment supported the services of two emulated CPs (CP A and CP B).

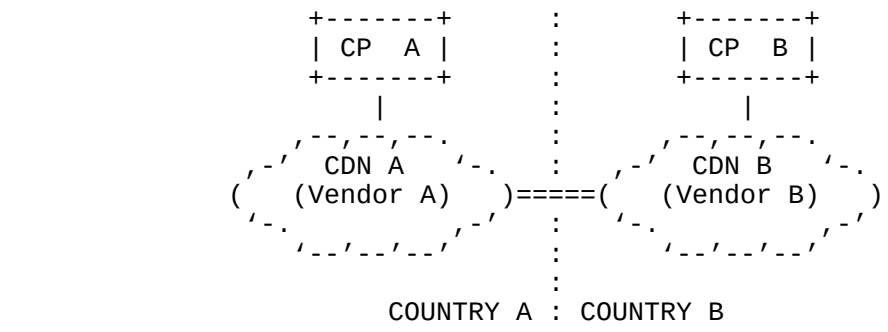

Figure 1

More precisely, we have run the experiments represented in Figure 2 and Figure 4. We base our description on Figure 2. In this experiment, CP A has an agreement with CDSP A for content delivery to end-users located in Country A and Country B. However, CDSP A operates a CDN (CDN A), whose footprint does not include country B. Therefore, CDSP A has an agreement with CDSP B, so that CDN A can delegate to CDN B the delivery of some content. More specifically, CDN A is allowed to delegate to CDN B the handling of requests sent by end-users located in Country B for CP A's content.

When CDN A receives a content request related to CP A and from an end-user in Country B, it redirects the end-user to CDN B. If CDN B does not have a local copy of the requested content yet (cache miss), CDN B ingests the content from CDN A (or from the CP's origin servers, depending on the test scenario); if CDN A also does not have a local copy of the requested content, it requests this asset from the CP's origin servers before sending the asset to CDN B.

There are several differences between the tests in Figure 2 and Figure 4, in addition to the different role played by the two CDN solutions. We list the main ones below.

- We have tested different content acquisition methods (see Section 2.6).

- Specific URL schemes were involved in providing content acquisition source information to the downstream CDN. As we have

tested different content acquisition methods, depending on which solution played the role of dCDN, the two solutions have used different URL schemes to address content. Therefore, the tests required the configuration of different content delivery metadata on the uCDN (see Section 2.4).

- o The two solutions use different methods to identify the end-user's geographic locations (see Section 2.4).

2.2. Control

The tested CDN solutions support control APIs but those are proprietary, so that the tested CDN solutions do not support a common inter-operable CDNI control interface. Therefore, we have not tested CDNI control operations and we had to perform manually most operations related to the configuration of the CDNI.

2.3. Logging

Proprietary mechanisms to export transaction logs [I-D.bertrand-cdni-logging] were available in the tested CDN solutions, but have not been covered by our tests.

2.4. Request Routing and Content Delivery

As defined in [I-D.davie-cdni-framework], two main types of request-routing call flows can be used for CDNI:

1. iterative request-routing,

2. recursive request-routing.

Moreover, two main methods can be used for redirecting an end-user from the authoritative CDN to the downstream CDN:

1. DNS-based redirection,

2. Service-level redirection, for example HTTP-based or RTSP-based.

We have focused our tests on iterative request-routing. Our tests involved HTTP-based request redirection by the authoritative CDN, as they focused on the delivery of large objects, such as movies.

The tested CDN solutions did not feature a CDNI request-routing interface allowing exchange of "CDN routing" information among CDNs. Therefore, we have manually configured appropriate policies on the authoritative CDN to permit iterative request-routing (e.g., CDN A redirects the end-user to CDN B request-routing system, cf.

Figure 3).

2.4.1. HTTP Redirection by CDN A and Delivery by CDN B

This section describes the tested request routing and content delivery features in the scenario depicted in Figure 2, with HTTP redirection by CDN A and delivery by CDN B.

We have tested the selection of the downstream CDN based on the content type in the client request and/or the geographic location of the client.

- File-type based selection relied on static XML files where content file extensions can be associated to a specific delivery CDN.
- The geographic location of the end-user was determined by an external geolocation server.

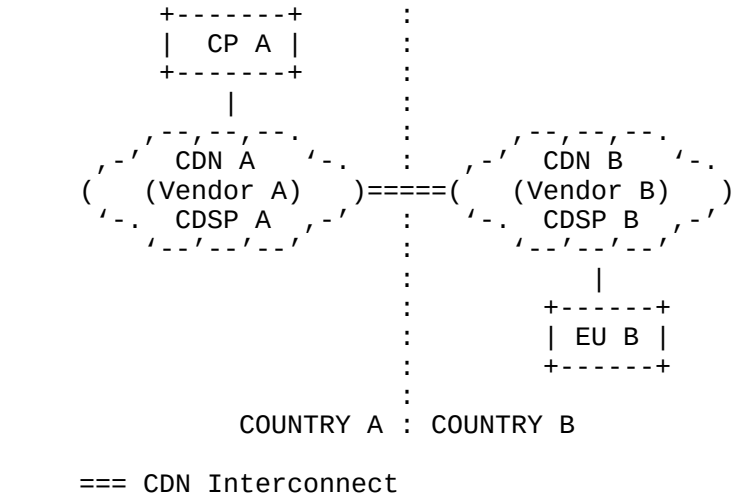

Figure 2

Figure 3 details the messages exchanged by the components involved in the experiment.

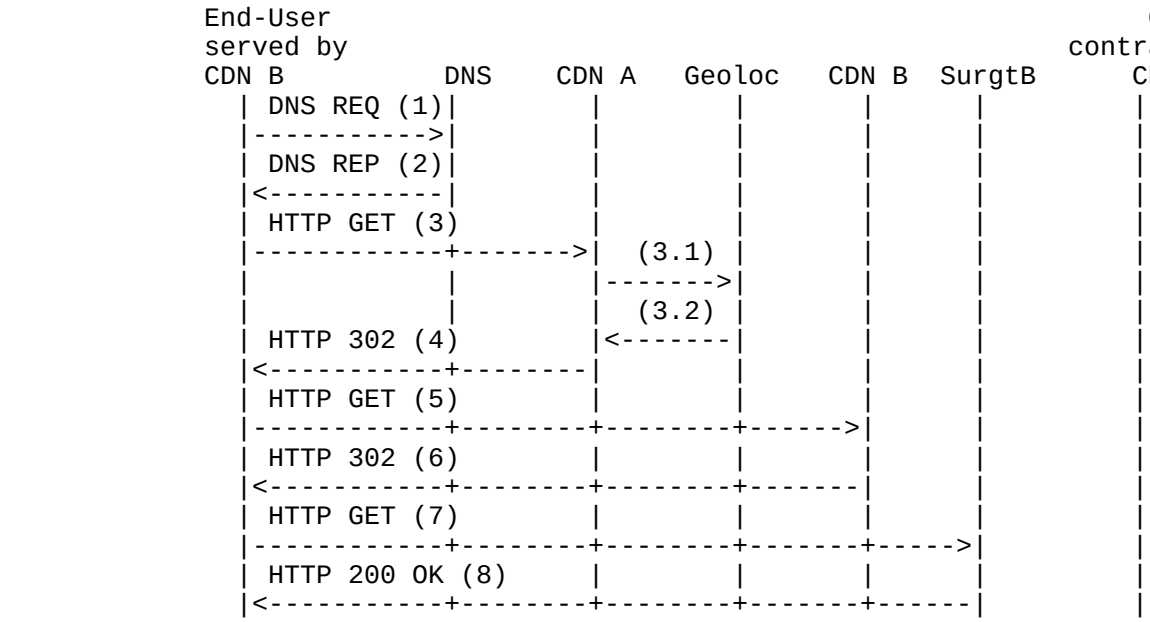

HTTP redirection by CDN A and delivery by CDN B

Figure 3

Message details

(1) The user-agent sends a DNS request to resolve the FQDN of the content URI.

(2) The DNS answers with the IP address of a request-router in CDN A.

(3) The user-agent sends to CDN A request-router an HTTP GET request for the content URI.

(3.1) CDN A request-router analyzes the request and queries a geolocation database to identify the geographic location of the end-user.

(3.2) The geolocation database answers with geolocation information related to the end-user’s IP address. The end-user is in country B; thus, CDN A determines that the end-user’s request must be served by CDN B.

(4) CDN A request-router replies to the user-agent with an HTTP 302 redirection message, which provides the URI of the content on CDN B.

(5) If necessary, the user-agent resolves the FQDN on the redirection URI (steps not represented in the figure), and thus, determines the IP address of a request-router in CDN B. Then, it sends an HTTP GET request to this request-router.

(6) CDN B request-router analyzes the request and replies to the user-agent with an HTTP 302 redirection message that provides the URI of the content on a surrogate in CDN B.

(7) If necessary, the user-agent resolves the FQDN of the redirection URI (steps not represented in the figure), and thus, determines the IP address of a surrogate in CDN B. Then, it sends an HTTP GET request to the surrogate.

(8) The surrogate analyzes the request and delivers the requested content to the end-user, through an HTTP 200 OK message.

2.4.2. HTTP Redirection by CDN B and Delivery by CDN A

This section describes the tested request routing and content delivery features in the scenario depicted in Figure 4, with HTTP redirection by CDN B and delivery by CDN A.

We have tested the selection of the downstream CDN based on end-user’s geolocation. For these tests, the geolocation database had been populated manually with the mapping of IP prefixes to countries. Alternative solutions, such as geolocation based on BGP communities or on the extraction of per country IP prefixes thanks to commercial geoIP databases, exist but they have not been tested in this experiment.

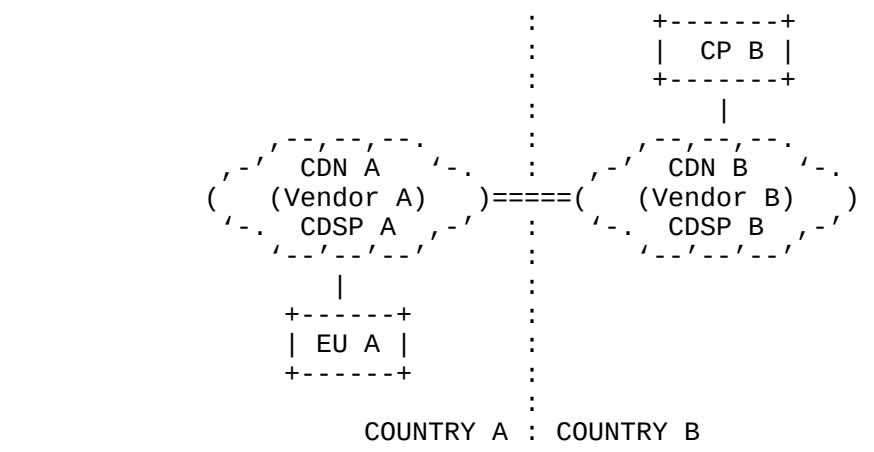

Figure 4

Figure 5 details the messages exchanged by the components involved in the experiment.

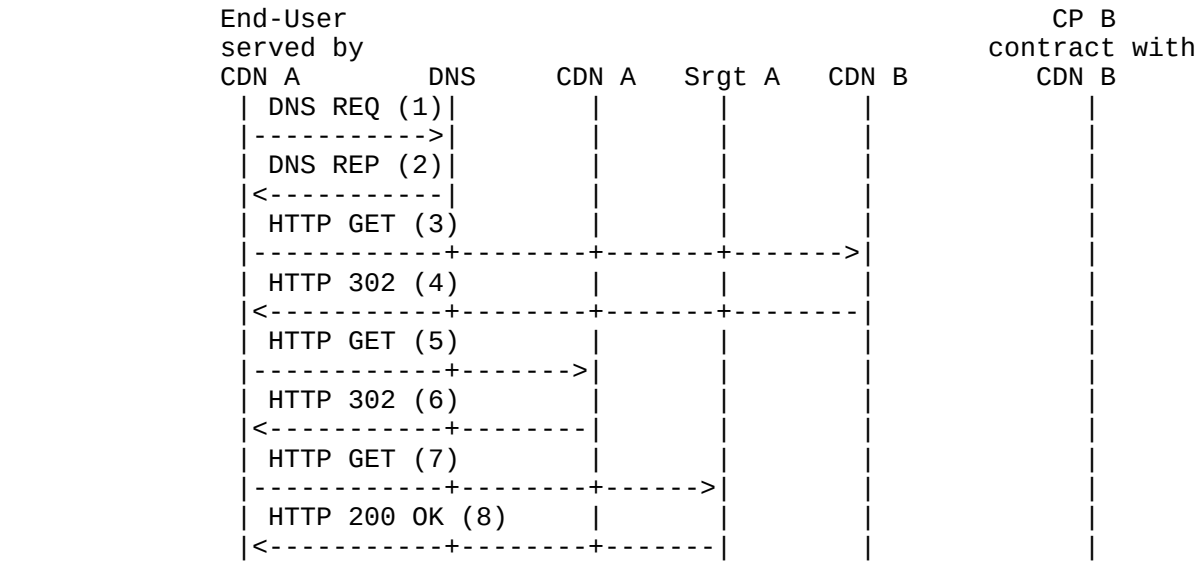

HTTP redirection by CDN B and delivery by CDN A

Figure 5

Message details

(1) The user-agent sends a DNS request to resolve the FQDN of the content URI.

(2) The DNS answers with the IP address of a request-router in CDN B.

(3) The user-agent sends to CDN B request-router an HTTP GET request for the content URI.

(4) CDN B request-router analyzes the request. The end-user is in country A; thus, CDN B determines that the end-user's request must be served by CDN A. Consequently, CDN B replies to the user-agent with an HTTP 302 redirection message that provides the URI of the content on CDN A.

(5) If necessary, the user-agent resolves the FQDN on the redirection URI (steps not represented in the figure), and thus, determines the IP address of a request-router in CDN A. Then, it sends an HTTP GET request to this request-router.

(6) CDN A request-router analyzes the request and replies to the user-agent with an HTTP 302 redirection message, which provides the URI of the content on a surrogate in CDN A.

(7) If necessary, the user-agent resolves the FQDN of the redirection URI (steps not represented in the figure), and thus, determines the IP address of a surrogate in CDN A. Then, it sends an HTTP GET request to the surrogate.

(8) The surrogate analyzes the request and delivers the requested content to the end-user, through an HTTP 200 OK message.

2.4.3. Test Result

HTTP redirection by the authoritative CDN was successful in the tests: end-users were redirected to the CDN that served their country. This guaranteed that:

- content from CP A be delivered by CDN B to end-users in country B, even if CP A had no direct relationship with CDSP B;
- content from CP B be delivered by CDN A to end-users in country A, even if CP B had no direct relationship with CDSP A.

2.5. Content Delivery Metadata

The tested CDN solutions feature proprietary metadata APIs, but these APIs have not been covered by the tests. We had to configure distribution metadata consistently in the dCDN and the uCDN (e.g., rules to determine upstream source for content acquisition).

Content pre-positioning in the dCDN has not been tested: only dynamic content acquisition has been covered by the experiments.

Proprietary APIs were available for content purge, but those have not been covered by tests.

2.6. Content Acquisition

We have used regular HTTP for content acquisition. We have relied on HTTP custom headers to transfer trivial metadata such as content integrity check (MD5 hash).

The correct acquisition and delivery of the requested file has been tested for Adobe Flash and MS HTTP smooth-streaming files.

2.6.1. Content Acquisition by CDN B through CDN A

We describe here the content acquisition operations triggered in case of cache miss, for the test scenario depicted in Figure 2 and Figure 3, with HTTP redirection by CDN A and delivery by CDN B. In this scenario, the dCDN (CDN B) does not have the requested content in cache and must request it to the uCDN (CDN A).

The uCDN treats the dCDN surrogate as an end-user: Figure 6 provides a summary (the involved internal entities of the uCDN are not detailed) of the related content acquisition operations. Section 3.1.3 provides more details on specific issues related to this content acquisition mode.

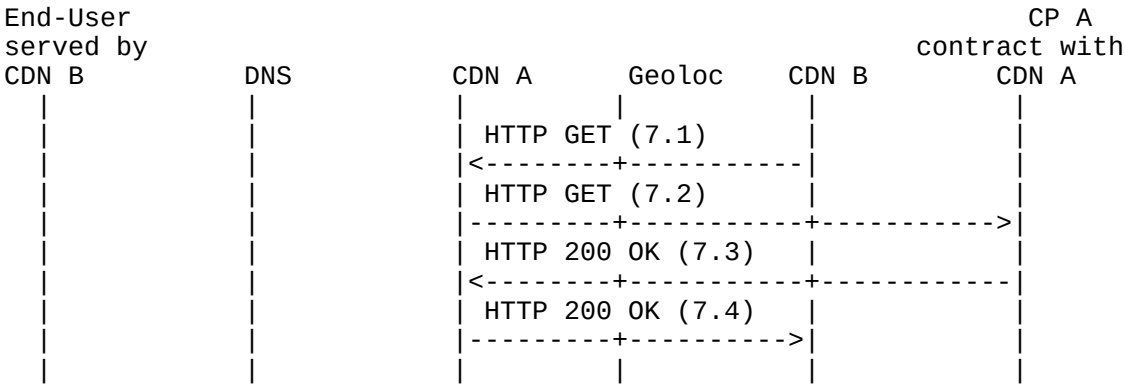

Pull content acquisition by CDN B through CDN A in case of cache miss
(continuation of Figure 3)

Figure 6

Message details

(7.1) CDN B surrogate or parent cache sends a content acquisition request to the uCDN (CDN A). In the tests, (7.1) was triggered by a cache miss on delivery request. Stated differently, the tests implemented dynamic acquisition, as opposed to content pre-positioning.

(7.2) CDN A analyzes the request. In case of cache miss, it sends an acquisition request to the CP's origin servers.

(7.3) The CP's origin server authorizes the request and delivers the requested content to CDN A, through an HTTP 200 OK.

(7.4) CDN A delivers the requested content to CDN B surrogate or parent cache, through an HTTP 200 OK.

2.6.2. Content Acquisition by CDN A Directly from CP B

We describe here (Figure 7) the content acquisition operations triggered in case of cache miss, for the test scenario with HTTP redirection by CDN B and delivery by CDN A (Figure 4 and Figure 5). In this scenario, the dCDN (CDN A) does not have the requested content in cache and must request it to CP B.

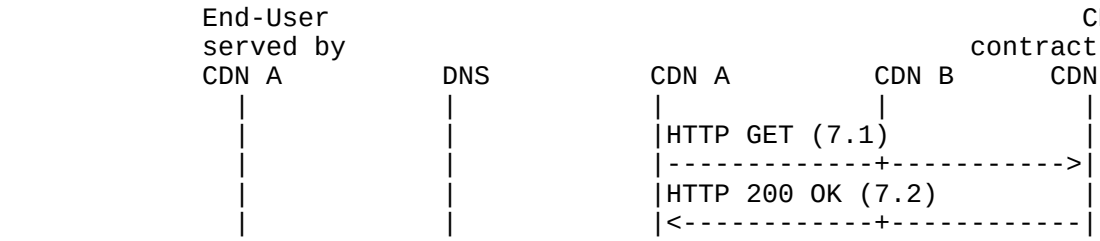

Pull content acquisition by CDN A directly from content provider's origin servers in case of cache miss (continuation of Figure 5)

Figure 7

Message details

(7.1) CDN A surrogate sends a content acquisition request to an origin server of the CP. In the tests, (7.1) was triggered by a cache miss on a delivery request. Stated differently, the tests implemented dynamic acquisition, as opposed to content pre-positioning.

(7.2) The CP's origin server authorizes the request and delivers the requested content to CDN A, through an HTTP 200 OK.

3. Lessons Learned

For basic interconnection tests, we have relied on extremely limited information exchanges between the two interconnected CDNs and we have configured most CDNI related features manually. This is because, while present CDN technologies support APIs allowing configuration of some of this information, those are difficult to use in multi-vendor environments since:

- they are proprietary APIs;
- they are designed as "internal" APIs and therefore lack the necessary inter-domain security and policy control.

Therefore, those APIs have not been used in these tests.

In the present section, we highlight some of the limitations induced by the lack of standard CDNI interfaces that we have faced in our tests.

One of the insights from this work is that by encoding information

inside the redirection URI, it is possible to communicate some essential CDNI-related information across CDNs "in-band" (i.e., as part of HTTP), rather than communicating it through an out-of-band interface. In these tests, the information communicated in-band was restricted to the most fundamental information; that is, a handle allowing the Downstream CDN to determine where to acquire the content. This was key to achieving multi-CDN operations without any common CDNI "out-of-band" interface supported by existing CDN technologies. This raises an interesting general question: what subset of inter-CDN information is to be communicated between interconnected CDNs in-band (possibly using existing methods) as opposed to communicated via out-of-band interfaces?

3.1. Request Routing

3.1.1. Request-Routing Information and Policies

Because of the lack of CDNI interfaces allowing CDNs to exchange information such as their coverage, capabilities, and performance, we had to configure request-routing policies manually in the CDNs that acted as uCDNs. While this may be tolerable for initial limited deployments of CDNI scenarios with a small number of participants, this is expected to create operational constraints in larger scale deployments.

3.1.2. Iterative and Recursive Redirection

Because of the lack of CDNI interfaces allowing an upstream CDN to query dCDN for how to redirect a request, the tests only covered iterative redirection (i.e. uCDN redirects the user-agent to the dCDN request-routing system, which redirects the user-agent to ...), not recursive redirection.

While iterative redirection allows supporting redirection across CDNs, it has some limitations:

- multiple redirections are exposed to the end-user;
- redirection latency cannot easily be reduced for future requests, through the caching of request-routing decisions;
- some client implementations support a limited number of successive redirections;
- the dCDN cannot reject a redirection, while allowing the uCDN to handle the rejected request.

A standard request-routing API would allow supporting recursive

redirection, which removes these shortcomings.

3.1.3. Request Looping Avoidance

In case of cache-miss, the downstream CDN must fetch the requested content, either through the authoritative CDN, or directly from the CP's origin server.

Consider the situation where the downstream CDN fetches the content from the authoritative CDN, as illustrated in Section 2.6.1. In this case, the authoritative CDN must not redirect the acquisition request to the downstream CDN, because this would create the following request-routing loop: dCDN -> uCDN -> dCDN. Consequently, the upstream CDN must be able to determine that the source of the request is a partner CDN and not a regular end-user. In addition, the upstream CDN must be able to acquire content from the CP's origin server on behalf of the downstream CDN, if necessary: dCDN -> uCDN -> CP.

In the tests, we have successfully solved the request-looping issue, through the use of separated URL spaces for regular users and CDN users, as well as the manual configuration of appropriate request-routing policies for every URL space. In other words, the URL used by regular users to fetch content was different from the one used by the downstream CDN to fetch the content. This way, we eliminated the loops. More automated operation would be required in larger-scale deployments.

3.2. Content Delivery Metadata

CDN technology typically supports APIs allowing creation, update, and deletion of content delivery metadata in the CDN. However, while often similar, those are proprietary and would require custom support. In the tests, passing of the most essential information, i.e., the upstream source for content acquisition, was achieved indirectly via conveying a handle inside the URI and configuring manually in the downstream CDN rules for extracting the upstream source.

The upstream source for content acquisition can be specified through the use of a specific URI scheme. For example, CDN A could use the following scheme: http://cdni.cdna.com/origin-URI to point to a cached copy of the content reachable at "origin-URI". The domain name inside the URI scheme designates the request-routing system of a CDN, and the remainder of the URI defines upstream source for content acquisition: here, the content URI on the CP's origin servers. If the authoritative CDN and the downstream CDN use this URI scheme, the authoritative CDN can easily map the URI that it receives in the end-

user's content request with a valid URI on the downstream CDN. Similarly, the downstream CDN can easily extract a content acquisition URI from the redirection URI.

In our tests, we have configured manually the rules that enabled the authoritative CDN to redirect end-users to a valid URI on the downstream CDN, and the downstream CDN to ingest content from the appropriate upstream source. While this allowed validation of the content distribution model, this solution may not be viable in a production environment, as it imposes constraints on URI structure. In addition, this approach does not support exchange of other distribution metadata (e.g., geo-blocking, content validation,...) which would require to be manually configured in the downstream CDN. In a real-world deployment, the configuration of these policies could rely on information that interconnected CDNs would exchange through a CDNI interface.

3.3. Content Acquisition and Deletion

3.3.1. Content Pre-Positioning in Downstream CDN

CDN technology typically supports APIs that allow triggering of content and metadata pre-positioning in a CDN. However, while often similar, these APIs are proprietary and would require custom support. For this reason content pre-positioning in dCDN was not covered in the tests. While the highest requirements is for support of dynamic acquisition, CDNI use-cases call for support of pre-positioning, which requires a triggering mechanism in a CDNI API.

3.3.2. Content Purge

CDN technology typically supports APIs allowing content purge in a CDN. However, while often similar, these APIs are proprietary and would require custom support. For this reason, content purge was not covered in the tests. There is a strong requirement for content purge in CDNI scenarios, which introduces the need for a purge triggering mechanism in a CDNI API.

4. Acknowledgments

The authors would like to acknowledge the work of Elodie Hemon and Marcin Pilarski on the tests, with the technical support of Sharon Schwartzman and Marc Fiuczynski. They would like to thank Slim Gara, Vincent Lauwers, Emile Stephan, Benoit Gaussen, and Mateusz Dzida for valuable input and discussions. Finally, the authors acknowledge interesting discussions with contributors of the EU FP7 OCEAN project.

5. IANA Considerations

This memo includes no request to IANA.

6. Security Considerations

CDN interconnect, as described in this document, has a wide variety of security issues that should be considered. This document focuses on specific experiments for CDN interconnect, and therefore, does not analyze the threats in detail.

7. References

7.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

7.2. Informative References

[I-D.bertrand-cdni-logging]
Gilles, B. and S. Emile, "CDNI Logging Interface", draft-bertrand-cdni-logging-00 (work in progress), February 2012.

[I-D.davie-cdni-framework]
Davie, B. and L. Peterson, "Framework for CDN Interconnection", draft-davie-cdni-framework-01 (work in progress), October 2011.

[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-03 (work in progress), January 2012.

[I-D.ietf-cdni-requirements]
Leung, K. and Y. Lee, "Content Distribution Network Interconnection (CDNI) Requirements", draft-ietf-cdni-requirements-02 (work in progress), December 2011.

[I-D.ietf-cdni-use-cases]
Gilles, B., Emile, S., Watson, G., Burbridge, T., Eardley, P., and K. Ma, "Use Cases for Content Delivery Network Interconnection", draft-ietf-cdni-use-cases-03 (work in

progress), January 2012.

[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.

Authors' Addresses

Gilles Bertrand (editor)
France Telecom - Orange
38-40 rue du General Leclerc
Issy les Moulineaux 92130
FR

Phone: +33 1 45 29 89 46
Email: gilles.bertrand@orange.com

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

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

Larry Peterson
Verivue, Inc.
2 Research Way
Princeton, NJ 08540
US

Phone: +1 978 303 8032
Email: lpeterson@verivue.com

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: January 8, 2012

G. Bertrand
E. Stephan
France Telecom - Orange
G. Watson
T. Burbridge
P. Eardley
BT
K. Ma
Azuki Systems

July 7, 2011

Use Cases for Content Delivery Network Interconnection
draft-bertrand-cdni-use-cases-02

Abstract

Content Delivery Networks (CDNs) are commonly used for improving the footprint and the end-user experience of a content delivery service, at a reasonable cost. This document outlines real world use-cases (not technical solutions) for interconnecting CDNs. It provides the business motivations for CDNI Working Group, which can be used to validate different interconnection arrangements, and requirements of the various CDNI interfaces.

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 January 8, 2012.

Copyright Notice

Table of Contents

1. Introduction

This document now merges input from [I-D.watson-cdni-use-cases] and [I-D.ma-cdni-publisher-use-cases].

Content Delivery Networks (CDNs) are commonly used for improving the footprint and the end-user experience of a content delivery service, at a reasonable cost. This document outlines real world use-cases (not technical solutions) for interconnecting CDNs. It provides the business motivations for CDNI Working Group, which can be used to validate different interconnection arrangements, and requirements of the various CDNI interfaces.

There are many possible combinations for the relationships between the different parties (Network Service Provider (NSP), CDN Provider, Content Service Provider (CSP) and End User) involved in end-to-end content delivery. However, in the context of interconnecting CDNs the key relationships are listed below.

- How the CSP interacts with the CDN provider, so that the CDN delivers content in a manner compliant with CSP's distribution policies.
- How the End User interacts with the CSP and one or more CDNs to request and receive content.
- How the different CDN providers, operating their CDNs, interact with one another to deliver the CSP's content to the End User while continuing to enforce the CSP's distribution policies.

This document describes a number of use cases that motivate CDN Interconnection.

1.1. Terminology

We adopt the terminology described in [I-D.jenkins-cdni-problem-statement], [RFC3466], and [RFC3568], except for the terms defined below.

CDN Provider:

An administrative entity who operates a CDN over a NSP or over the Internet.

Authoritative CDN (aCDN):

A CDN provider contracted by the CSP for delivery of content by its CDN or by its downstream CDNs.

Downstream CDN (dCDN):

A CDN provider which is contracted by an uCDN to achieve the delivery of content to users.

Access CDN:

A CDN that is connected to the end-user's access and has information about the end-user's profile and access capabilities.

Delivering CDN:

The CDN that delivers the requested content asset to the end-user. In particular, the delivering CDN can be an access CDN.

CDN Interconnection (CDNI):

Relationship between two CDNs that enables a CDN to provide content delivery services on behalf of another CDN. It relies on a set of interfaces over which two CDNs communicate in order to achieve the delivery of content to end-users by one CDN (the downstream CDN) on behalf of another CDN (the upstream CDN).

CDN peering: A business relation between two CDN providers based on one or more CDN interconnections.

Recursive request routing:

Recursive: Where a process is repeated, but embedded within the original process. In the case of Request Routing, this means that the initial request received by the Authoritative CDN is processed downstream from one CDN to another and that the responses are send back upstream to the Authoritative CDN which then replies to the initial request.

Iterative request routing

Iterative: Where a process is repeated multiple times to make progress towards a goal. In the case of Request Routing, this means that the initial request is received by the Authoritative CDN, which replies it with a redirection directive to a downstream CDN. When the end-user sends its request to the downstream CDN, the same process is repeated, until the request arrives to the delivering CDN.

Asymmetric Distribution:

A distribution scenario where different NSPs have distribution rights to the same content, but at different levels of quality (e.g., high

definition vs. low definition video), which places restrictions on delivery delegation.

1.2. Abbreviations

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

- CSP: Content Service Provider
- dCDN: downstream CDN
- ISP: Internet Service Provider
- NSP: Network Service Provider
- PC: Personal Computer
- QoE: Quality of Experience
- QoS: Quality of Service
- SLA: Service Level Agreement
- STB: Set-Top-Box
- uCDN: upstream CDN
- UA: User Agent
- UE: User Equipment
- VoD: Video on Demand
- WiFi: Wireless Fidelity

1.3. High Level Use Cases for Multi-CDN Systems

Content Delivery Networks (CDNs) are used to deliver content because they can:

- improve the experience for the End User; for instance delivery has lower latency and better robustness,
- reduce the operator’s costs; for instance lower delivery cost (reduced bandwidth usage) for cacheable content,
- reduce the Content Service Provider costs, such as datacenter capacity, space, and electricity consumption.

Indeed, many network service providers and enterprise service providers are deploying or have deployed their own CDNs. Despite the potential benefits of interconnecting CDNs, today each CDN is a standalone network. The objective of CDN interconnection is to overcome this restriction: the interconnected CDNs should be able to collectively behave as a single delivery infrastructure.

Let's take an example, as depicted in Figure 1. Two CDN Providers establish a CDN Interconnect. The Content Service Provider CSP-1 reaches an agreement with CDN Provider A for the delivery of its content. CDN Provider A and CDN Provider B agree to interconnect their CDNs. When a User Agent that is connected to CDN Provider B's network requests Content from CSP-1, the Content is actually delivered from CDN-B, because handling of requests for CSP-1's Content has been delegated as part of the CDN Interconnect agreement. The End User benefits through a better quality of experience, because the Content is delivered from a nearby Surrogate. CDN Provider A benefits because it doesn't need to deploy such an extensive CDN, whilst CDN Provider B receives some compensation for the delivery. CSP-1 benefits because it only needs to make one business agreement and one physical connection, with CDN Provider A, but its End Users get a service quality as though it had also gone to the trouble of making a business agreement with CDN Provider B.

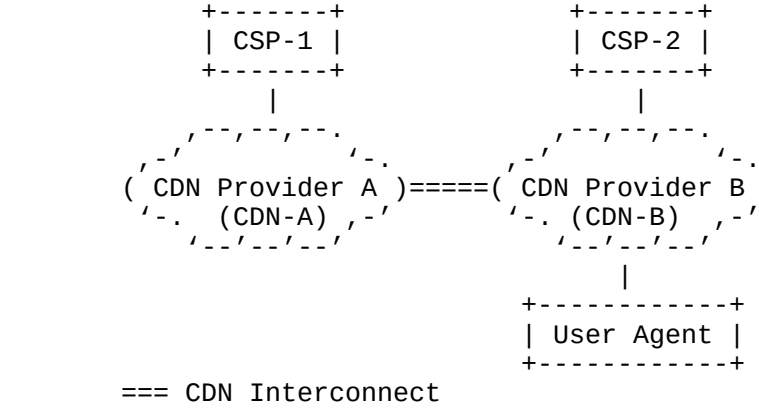

Figure 1

To extend the example, another Content Service Provider, CSP-3, may also reach an agreement with CDN Provider A. But it does not want its Content to be distributed by CDN Provider B, for example, CSP-3 may not have distribution rights in the country where CDN Provider B operates. This example illustrates that policy considerations are an

important part of CDNI.

This document identifies three main motivations for a CDN Provider to interconnect its CDN:

- CDN Footprint Extension Use Cases (Section 2)
- CDN Offload Use Cases (Section 3)
- CDN Capability Use Cases (Section 4)

1.4. The Need for CDNI Standards

Existing CDN interfaces are proprietary and an external CDN typically cannot use them, especially if the two CDNs rely on different solutions. Nevertheless, [I-D.bertrand-cdni-experiments] shows that some level of CDN interconnection can be achieved experimentally without standardized interfaces between the CDNs. The methods used in these experiments are hardly usable in an operational context, because they suffer from several limitations in terms of functionalities, scalability, and security level.

The aim of the CDNI standards work is therefore to overcome such shortcomings; a full list of requirements is being developed in [I-D.lefaucheur-cdni-requirements].

2. Footprint Extension Use Cases

Footprint extension is expected to be a major use case for CDN interconnection.

2.1. Geographic Extension

In this use case, the CDN Provider wants to extend the geographic distribution that it can offer CSPs, without

- compromising the quality of delivery
- attracting transit and other network costs by serving from geographically or topologically remote surrogates.

If there are several CDN Providers that have a geographically limited footprint (e.g., restricted to one country), or do not serve all end-users in a geographic area, then interconnecting their CDNs enables CDN Providers to provide their services beyond their own footprint.

As an example, suppose a French CSP wants to distribute its TV

programs to End Users located in various countries in Europe and North Africa. It asks a French CDN Provider to deliver the content. The French CDN Provider's network only covers France, so it makes an agreement with another CDN Provider that covers North Africa. Overall, from the CSP's perspective the French CDN Provider provides a CDN service for both France and North Africa.

In addition to video, this use case applies to other types of content such as automatic software updates (browser updates, operating system patches, virus database update, etc).

2.2. Region to Region Interconnection

In the previous section, we have described the case of geographic extension between CDNs operated by different entities. A large CDN Provider may also operate CDNs from several subsidiaries (which may rely on different CDN solutions, see Section 4.2). In certain circumstances, the CDN Provider needs to make its CDNs interoperate to provide a consistent service to its customers on its whole footprint.

2.3. Nomadic Users

In this scenario a CSP wishes to allow users who move to other geographic regions to continue to access their content. The motivation in this case is to allow nomadic users to maintain access, rather than to allow all residents within a region access to the content.

This use case covers situations like users moving between different CDN Providers within the same geographic region, or users switching between different devices, as discussed in Section 4.

2.4. Delivery Restrictions

The content distribution policies that a CSP attaches to a content asset depend on many criteria. Distribution rights for audiovisual content are often negotiated using a combination of temporal licensing (e.g., available for 24 hours, available 28 days after DVD release, etc.), resolution-based licensing (e.g., high definition vs. standard definition), and geo- location-based licensing (e.g., per country).

"Geo-blocking" rules may specify:

- the geographic regions where content can be delivered from (i.e. the location of the Surrogates), or

- o geographic locations where content can be delivered to (i.e., the location of the End Users).

Hence, the exchange through the CDN interconnection of information for controlling the footprint of the delivery is an important use case.

The delivery of content may be further influenced by policies which may include time-based rules that specify:

- o an activation time (i.e., the time when the content should become available for delivery),
- o a deactivation time (i.e., time after which the content should no longer be delivered), or
- o an expiration time (i.e., the time at which the content files should be expunged from all CDN storage).

The delivery of content may be further influenced by policies which may include quality of service rules that specify:

- o the maximum resolution deliverable to specific devices,
- o the maximum resolution deliverable though a specific NSP, or
- o the maximum resolution deliverable to users based on their subscription levels.

The enforcement of CSP licensing rules when making CDN delegation decisions is another important use case for CDN interconnection.

3. Offload Use Cases

3.1. Overload Handling and Dimensioning

A CDN is likely to be dimensioned to support the prime-time traffic. However, unexpected spikes in content popularity may drive load beyond the expected peak. The prime recurrent time peaks of content distribution may differ between two CDNs. Taking advantage of the different traffic peak times, a CDN may interconnect with another CDN to increase its effective capacity during the peak of traffic. This brings dimensioning savings to the CDNs as they can use the resources of each other during their peaks of activity.

Offload also applies to planned situations where a CDN Provider needs CDN capacities in a particular region during a short period of time.

For example, a CDN can offload traffic to another CDN during a specific maintenance operation or for covering the distribution of a special event. For instance, consider a TV-channel which has exclusive distribution rights on a major event, such as a celebrities' wedding, or a major sport competitions. The CDNs that the TV-channel uses for delivering the content related to this event are likely to experience a flash crowd during the event and to need offloading traffic, while other CDNs will support a more usual traffic load and be able to handle the offloaded traffic load.

3.2. Resiliency

3.2.1. Failure of Content Delivery Resources

It is important for CDNs to be able to guarantee service continuity during partial failures (e.g., failure of some Surrogates). In partial failure scenarios, a CDN Provider could redirect some requests towards another CDN, which must be able to serve the redirected requests or, depending on traffic management policies, to forward these requests to the CSP's origin server.

3.2.2. Failure of Content Acquisition

Source content acquisition is typically handled in one of two ways:

- CDN origin, where a downstream CDN acquires content from an upstream CDN, and the authoritative CDN acquires content from an origin server of the CSP, or

- CSP origin, where the CDNs acquire content directly from an origin server of the CSP.

Resiliency may be required against failure to ingest content from the CSP. If a CDN is unable to retrieve the content, it may be that the CSP's origin server is inaccessible to only this CDN, in which case redirection of the end-users to an alternative CDN may circumvent the problem. A CSP may also choose to specify one or more backup origin servers.

3.3. Branding Consideration

There are situations where one CDN Provider cannot or does not want to operate all the functions of a CDN. For instance, it always acts as an uCDN and offloads the content delivery to dCDNs, i.e., it uses the surrogates of other CDSPs. In this model, the uCDN acquires content and receives the initial routing requests from the user agent; whereas, the dCDNs operate the content delivery functions. The uCDN also retrieves and presents the logging for the CSP.

Preserving branding elements could interest the CSP or CDSPs. The CSP might desire to offer content services under its name, even if the associated CDN service involves other organizations. Therefore, the CSP could request that the name of the CDSPs does not appear in the URLs. Similarly, in offload situations, the uCDN might want to offer CDN services under its own branding. This highlight a requirement for exchanging branding related constraints over a CDNI.

4. CDN Capability Use Cases

4.1. Device and Network Technology Extension

In this use case, the CDN Provider may have the right geographic footprint, but wishes to support the delivery of content to alternative devices, such as smartphones connected to a mobile network. In this case, the CDN Provider may federate with another CDN Provider that offers service to these devices.

Consider the scenario shown in Figure 2. In this example, a nomadic user switches from a TV going through a cable provider to a smartphone going through a mobile operator. The CDN Provider on the cable network may wish to delegate delivery of Content to the CDN Provider on the mobile network. There are several possible differences that may arise in this use case compared with the ones discussed earlier, for example:

- the phone may require the Content at lower resolution than the TV;
- the CSP may want to license only lower resolution Content to CDN Provider 2;
- the CSP may not want CDN Provider 2 to deliver Content if the connection quality is below some threshold;
- the CSP may want to tailor the Content in some special way depending on whether the End User is on cable or mobile, for example, different adverts / DRMs / codecs / container formats / delivery protocols...

These examples suggest the requirement for Asymmetric Distribution of Content across the CDN interconnect. In the nomadic scenario, the switch of CDN should be as seamless as possible from the End User’s perspective.

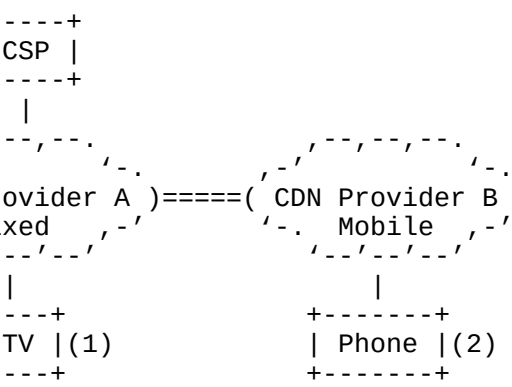

=== CDN Interconnect

Fixed-Mobile Session Shifting

Figure 2

4.2. Technology and Vendor Interoperability

A CDN Provider may deploy a new CDN to run alongside its existing CDN, as a simple way of migrating its CDN service to a new technology. A CDN Provider may have a multi-vendor strategy for its CDN deployment. A CDN Provider may want to deploy a separate CDN for a particular CSP or a specific network. In all these circumstances, CDNI benefits the CDN Provider, as it simplifies or automates some inter-CDN operations (e.g., migrating the request routing function progressively).

4.3. QoE and QoS Improvement

Some CSPs are willing to pay a premium for enhanced delivery of Content to their End Users. In some cases, even if the CDN Provider could deliver the content to the end users, it cannot meet the CSP's service level agreement. So, it makes a CDN Interconnect agreement with another CDN Provider that can meet the SLA, for instance an Access CDN, which is able to deliver content from Surrogates located closer to the end-user.

5. Acknowledgments

The authors would like to thank Francois Le Faucheur and Ben Niven-Jenkins for lively discussions.

They also thank the contributors of the EU FP7 OCEAN and ETICS projects for valuable inputs.

6. IANA Considerations

This memo includes no request to IANA.

7. Security Considerations

CDN interconnect, as described in this document, has a wide variety of security issues that should be considered. The security issues fall into three general categories:

- CSP Trust: where the CSP may have negotiated service level agreements for delivery quality of service with the uCDN, and/or configured distribution policies (e.g., geo-restrictions, availability windows, or other licensing restrictions), which it assumes will be upheld by dCDNs to which the uCDN delegates requests. Furthermore, billing and accounting information must be aggregated from dCDNs with which the CSP may have no direct business relationship. These situtations where trust is delegated must be handled in a secure fashion to ensure CSP confidence in the CDN interconnection.

- Client Transparency: where the client device or application which connects to the CDN must be able to interact with any dCDN using its existing security and DRM protocols (e.g., cookies, certificate-based authentication, custom DRM protocols, URL signing algorithms, etc.) in a transparent fashion.

- CDN Infrastructure Protection: where the dCDNs must be able to identify and validate delegated requests, in order to prevent unauthorized use of the network and to be able to properly bill for delivered content. A dCDN may not wish to advertise that it has access to or is carrying content for the uCDN or CSP, especially if that information may be used to enhance denial of service attacks. In general, CDNI interfaces and protocols should minimize overhead for dCDNs.

This document focuses on the motivational use cases for CDN interconnect, and does not analyze these threats in detail.

8. References

8.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

8.2. Informative References

[I-D.bertrand-cdni-experiments]
Bertrand, G., Faucheur, F., and L. Peterson, "Content Distribution Network Interconnection (CDNI) Experiments", draft-bertrand-cdni-experiments-00 (work in progress), February 2011.

[I-D.jenkins-cdni-problem-statement]
Niven-Jenkins, B., Faucheur, F., and N. Bitar, "Content Distribution Network Interconnection (CDNI) Problem Statement", draft-jenkins-cdni-problem-statement-02 (work in progress), March 2011.

[I-D.lefaucheur-cdni-requirements]
Faucheur, F., Viveganandhan, M., Watson, G., and Y. Lee, "Content Distribution Network Interconnection (CDNI) Requirements", draft-lefaucheur-cdni-requirements-01 (work in progress), March 2011.

[I-D.ma-cdni-publisher-use-cases]
Nair, R. and K. Ma, "Content Distribution Network Interconnection (CDNI) Publisher Use", draft-ma-cdni-publisher-use-cases-00 (work in progress), March 2011.

[I-D.watson-cdni-use-cases]
Watson, G., "CDN Interconnect Use Cases", draft-watson-cdni-use-cases-00 (work in progress), January 2011.

[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.

Authors' Addresses

Gilles Bertrand
France Telecom - Orange
38-40 rue du General Leclerc
Issy les Moulineaux, 92130
FR

Phone: +33 1 45 29 89 46
Email: gilles.bertrand@orange-ftgroup.com

Stephan Emile
France Telecom - Orange
2 avenue Pierre Marzin
Lannion F-22307
France

Email: emile.stephan@orange-ftgroup.com

Grant Watson
BT
pp GDC 1 PP14, Orion Building, Adastral Park, Martlesham
Ipswich, IP5 3RE
UK

Email: grant.watson@bt.com

Trevor Burbridge
BT
B54 Room 70, Adastral Park, Martlesham
Ipswich, IP5 3RE
UK

Email: trevor.burbridge@bt.com

Philip Eardley
BT
B54 Room 77, Adastral Park, Martlesham
Ipswich, IP5 3RE
UK

Email: philip.eardley@bt.com

Kevin Ma
Azuki Systems
43 Nagog Park
Acton, MA 01720
USA

Phone: +1 978 844 5100
Email: kevin.ma@azukisystems.com

Network Working Group
Internet-Draft
Intended status: Informational
Expires: September 14, 2011

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

March 13, 2011

Content Distribution Network Interconnection (CDNI) Problem Statement
draft-jenkins-cdni-problem-statement-02

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 creates a requirement 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 for the IETF with a view towards creating a working group. This working group would work on interoperable and scalable solutions for CDN interconnection.

Requirements Language

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 RFC 2119 [RFC2119].

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 September 14, 2011.

Table of Contents

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 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 End User regardless of that End User's location or attachment network. However, the footprint of a given CDN in charge of delivering a given content may not expand close enough to the End User's current location or attachment network to realize the cost benefit and user experience that a more distributed CDN would provide. This creates a requirement for interconnecting standalone CDNs so that their collective CDN footprint can be leveraged 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 for the IETF with a view towards creating a working group. This working group would work on interoperable and scalable solutions for CDN interconnection.

Section 2 discusses the use cases for CDN interconnection. Section 3 presents the CDNI model and problem area to be considered by the IETF. Section 4 discusses how existing protocols can be reused to define the CDNI protocols while Section 5 proposes to focus the scope for the initial charter of a CDNI Working Group to the minimum functional elements necessary for basic CDN interconnection. Section 5 provides a gap analysis of the work of other standards organization and finally Section 5 discusses the relationship with relevant IETF Working 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 (and not directly relevant to the distribution of that Content) or content representation. 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 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 Interconnect, dynamic acquisition means that a downstream CDN does not acquire the content from content sources (including upstream CDNs) until a request for that content has been delegated to the downstream CDN by an Upstream CDN.

Dynamic CDNI metadata acquisition: In the context of CDN Interconnect, dynamic CDNI metadata acquisition means that a downstream CDN does not acquire CDNI metadata for content from the

upstream CDN until a request for that content has been delegated to the downstream CDN by an Upstream CDN.

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

Pre-positioned CDNI Metadata acquisition: In the context of CDN Interconnect, Metadata Pre-positioning is where the the Downstream CDN acquires distribution 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 the Content Service. The User Agent will communicate with the CSP's 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 (STB), Dedicated content applications (e.g. media players), etc.

Network Service Provider (NSP): Provides network-based connectivity/ services to 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 the delivery of 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.

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. 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 Interconnect (CDNI): The set of interfaces over which two or more 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).

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

Downstream CDN: For a given 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. a CDN, operated by a different operator than the NSP to which the users of that service are attached.

Surrogate: A device/function 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.

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 Interconnect, 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 content 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 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.

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:

- RFC 3040 [RFC3040] describes many of the component technologies that are used in the construction of a CDN
- Taxonomy [TAXONOMY] compares the architecture of a number of CDNs
- 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 Interconnect 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 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. 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 number of End Users and/or over many/all geographies and/or with 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. 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, for this requirement, CDN operators (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. 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). Again, for this requirement, CDN operators (over-the-top CDN operators or NSPs) are faced with a lack of open specifications and best practices.

Use cases for CDN Interconnection are further discussed in [I-D.bertrand-cdni-use-cases] (which contains a merged set of use cases previously presented in [I-D.watson-cdni-use-cases] and [I-D.bertrand-cdni-use-cases-00]).

3. CDN Interconnect Model & Problem Area for IETF

Interconnecting CDNs involves interactions among multiple different functions and components that form each CDN. Only some of those require standardization. The CDNI model and problem area proposed for IETF work is illustrated in Figure 1. The candidate problem area (and respectively the non-goals) for IETF work on CDN Interconnection are discussed in Section 3.1 (and respectively Section 3.2).

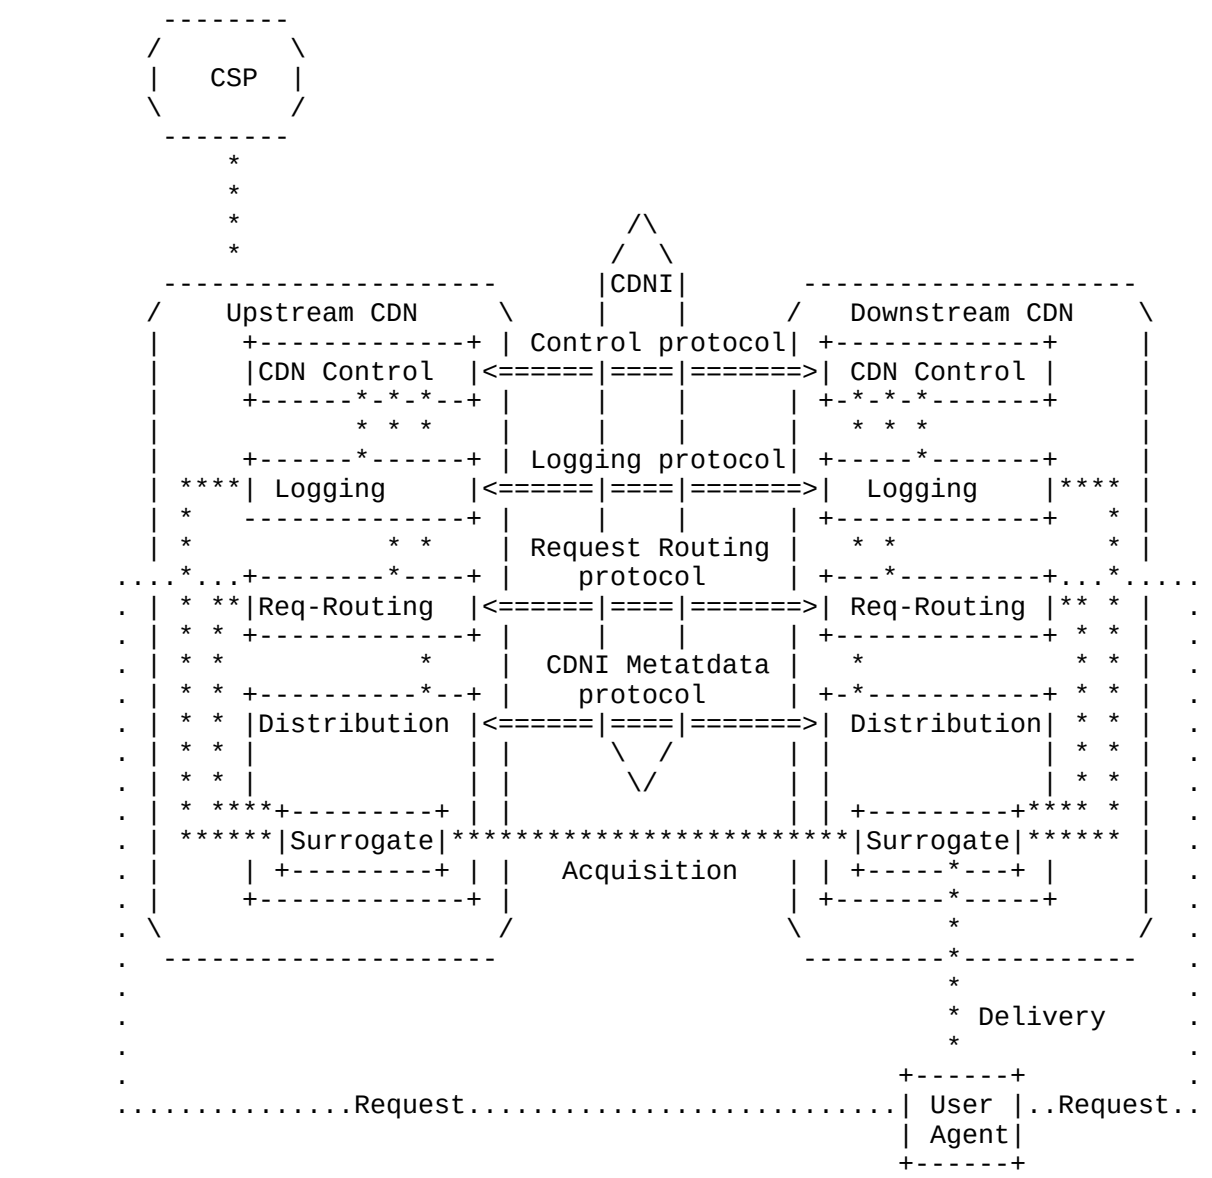

<==> interfaces inside the scope of CDNI

**** interfaces outside the scope of CDNI
.... interfaces outside the scope of CDNI

Figure 1: CDNI Problem Area

3.1. Candidate CDNI Problem Area for IETF

Listed below are the four protocols required to interconnect a pair of CDNs and that constitute the problem space that is proposed to be addressed by a potential CDNI working group in the IETF. The use of the term "protocol" is meant to encompass the protocol over which CDNI data representations (e.g. CDNI Metadata records) are exchanged as well as the specification of the data representations themselves (i.e. what properties/fields each record contains, its structure, etc.). While "interface" may be a more accurate term, the term "protocol" is retained in this document because of its common use.

- o CDNI Control protocol: This protocol allows the "CDNI Control" system in interconnected CDNs to communicate. This protocol may support the following:
  - * Allow bootstrapping of the other CDNI protocols (e.g. protocol address discovery and establishment of security associations).
  - * Allow configuration of the other CDNI protocols (e.g. Upstream CDN specifies information to be reported through the CDNI Logging protocol).
  - * 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 upstream CDN to initiate or request specific actions to be undertaken in the downstream CDN. For example, this may include the following capabilities:
    - + Allow an upstream CDN to request that content files and/or CDNI Metadata that it shared, be purged from, or invalidated in, a downstream CDN. Support for content deletion or invalidation from a CDN is a key requirement for some Content Service Providers in order, amongst other use cases for content deletion, to support the content rights agreements they have negotiated. Today’s CDNs use proprietary control interfaces to enable CSPs to remove content cached in the CDN and therefore there is a need to have a similar but standardized content deletion capability between interconnected CDNs.
    - + Allow an upstream CDN to initiate Pre-postioned content acquisition and/or Pre-positioned CDN Metadata acquisition in a downstream CDN.
- o CDNI Request Routing protocol: This protocol allows the Request Routing system 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 CDN). In particular, the CDN Request Routing protocol, may support the following:

* allow the upstream CDN to query the downstream CDN at request-routing time before redirecting the request to the downstream CDN
* allow 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 protocol: This protocol 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 protocol: This protocol 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 protocols is considered tentative at this stage and may be changed after further study (e.g. some subset of functionality be moved from one protocol 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 protocols need not be defined from scratch and instead can very significantly reuse or leverage existing protocols: this is discussed further in Section 4. Also, it is expected that the items above will be prioritized so that the CDNI Working Group can focus (at least initially) on the most essential and urgent work: this is discussed further in Section 5.

As part of the development of the CDNI protocols and solutions 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, as well as how to describe which policy should be used when doing so. [I-D.jenkins-cdni-names] presents one view on how CDN data types/objects could be classified such that the problem space of their naming and referencing is not as large as it

might at first appear because there is significant commonality between the different data types/objects required for CDNI.

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/protocols such as those discussed in this document is preventing the deployment of production CDN Interconnection solutions with the necessary level of functionality.

3.2. Non-Goals for IETF

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

- 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).
- The delivery interface between the delivering CDN surrogate and the User Agent, such as streaming protocols.
- 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 protocol may indirectly affect some of the information exchanged through the request interface (e.g. URI).
- 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 solution may 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.
- End User/User Agent Authentication. End User/User Agent authentication and authorization are the responsibility of the Content Service Provider.
- 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 Protocols. i.e. protocols within one CDN.
- o Scalability of individual CDNs. While scalability of the CDNI protocols/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.

The third bullet in the list above places the acquisition of content between interconnected CDNs as out of scope for CDNI and deserves some additional explanation. The consequence of such a decision is that a CDNI WG would be focussed on only defining the control plane for CDNI; and the CDNI data plane (i.e. the acquisition & distribution of the actual content objects) would not be addressed by a CDNI WG. 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 a CDNI WG is describing within the CDNI Metadata where to go and which protocol to use to retrieve the content.

4. Design Approach for Realizing the CDNI APIs

This section expands on how CDNI protocols can reuse and leverage existing protocols. First the "reuse instead of reinvent" design principle is restated, then each protocol is discussed individually with example candidate protocols that can be considered for reuse or leverage. This discussion is not intended to pre-empt any WG

decision as to the most appropriate protocols, technologies and solutions to select to solve CDNI but is intended as an illustration of the fact that these protocols need not be created in a vacuum and that reuse or leverage of existing protocols is likely possible.

4.1. Relationship to the OSI network model

The four CDNI protocols (CDNI Control protocol, CDNI Request Routing protocol, CDNI Metadata protocol, CDNI Logging protocol) described in Section 3.1 within the CDNI problem area are all control plane interfaces operating at the application layer (Layer 7 in the OSI network model). Since it is not expected that these protocols 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 protocols will be defined on top of existing session, transport and network protocols.

4.2. "Reuse Instead of Reinvent" Principle

Although a new application protocol could be designed specifically for CDNI we assume that this is unnecessary and it is recommended that existing application protocols be reused or leveraged (HTTP [RFC2616], Atom Publishing Protocol [RFC5023], XMPP [RFC3920], for example) to realize the CDNI protocols.

4.3. CDNI Request Routing API

The CDNI Request Routing protocol 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.

The CDNI Request Routing protocol needs to offer a mechanism for an upstream CDN to issue a "Redirection Request" to a downstream CDN. The Request Routing protocol 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 needs to contain information such as:

- The protocol (e.g. DNS, HTTP) over which the upstream CDN received the initial User Agent request
- 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 would need to contain 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 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 protocol may be cacheable.
- o Some CDNs may want broader policies, e.g. CDN B agrees to always take CDN A's delegated redirection requests, in which case the necessary redirection details are exchanged out of band (of the CDNI protocols), 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 needs 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 application specific redirection response(s) to return to the requesting User Agent. For HTTP requests from User Agents this could be in the form of a URI that the upstream CDN could return in a HTTP 302 response.

The CDNI Request Routing protocol is therefore a fairly straightforward request/response protocol 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 a CDNI WG to define a new protocol for the

request/response element of the Request Routing protocol. Thus, a CDNI WG would be left only with the task of specifying:

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

4.4. CDNI Metadata API

The CDNI Metadata protocol enables the Metadata function in a downstream CDN to obtain 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 protocol.
- Content Requests received directly from User Agents.

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

- distribute/update/remove CDNI Metadata to a Downstream CDN

and/or to allow a downstream CDN to:

- Make direct requests for CDNI Metadata records where the downstream CDN knows the identity of the Metadata record(s) it requires.
- Search for CDNI Metadata records where the downstream CDN does not know the specific Metadata record(s) it requires but does know some property of the record it is searching for. For example, it may know the value of the HTTP Host header received in a HTTP request and it wants to obtain the CDNI Metadata for that host so that it can determine how to further process the received HTTP request.

The CDNI Metadata protocol is therefore similar to the CDNI Request Routing protocol because it is a request/response protocol 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 protocol, the CDNI Metadata protocol 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 [RFC3920]. This removes the need for a CDNI WG to define a new protocol for the

request/response element of the Metadata protocol.

Thus, a CDNI WG would be left only with the task of specifying:

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

4.5. CDNI Logging API

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

- Log lines related to the delivery of content (similar to the log lines recorded in a web server's access log).
- Real-time or near-real time events before, during or after content delivery, e.g. content Start/Pause/Stop events, etc.
- 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 Operator 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 use cases place requirements on the CDNI Logging protocol 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 protocol, a CDNI WG would still need to

specify:

- The recommended protocol to use.
- 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.
- A default set of events that trigger logs to be generated.

4.6. CDNI Control API

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

However, as discussed in Section 3.1, the CDNI Control protocol may be required to support functionality similar to the following:

- Allow an upstream CDN and downstream CDN to establish, update or terminate their CDNI interconnection.
- Allow bootstrapping of the other CDNI protocols (e.g. protocol address discovery and establishment of security associations).
- Allow configuration of the other CDNI protocols (e.g. Upstream CDN specifies information to be reported through the CDNI Logging protocol).
- Allow the downstream CDN to communicate information about its delivery capabilities, resources 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).

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

5. Prioritizing the CDNI Work

In order to manage the potential workload of a CDNI WG, it is recommended that the work be prioritized in a "walk before you run" approach.

The CDNI problem area can be categorized into different solution scopes as follows:

- "Base CDNI" Scope: This solution scope comprises the solution elements that can be considered as the ’minimum’ needed to actually deliver any content using interconnected CDNs. For

example, a base CDNI Request Routing protocol and a base CDNI Metadata protocol belong to this scope because without them the upstream CDN is unable to redirect User Agents to the downstream CDN and the downstream CDN is unable to obtain the delivery policies and other CDNI Metadata required to ingest and deliver the content.

- o "Operationalized CDNI" Scope: This solution scope comprises the solution elements that can be considered as the 'minimum' needed to 'operationalize' CDN Interconnects. For example, the CDNI Logging protocol and the base capabilities of the CDNI Control protocol (e.g. content file/metadata deletion) belong to this scope because without them CDN operators are required to substitute for them either with manual processes or proprietary interfaces.
- o "Enhanced CDNI" Scope: This solution scope comprises the solution elements that can be classed as 'enhanced features'. For example, the aspects of the CDNI Control protocol related to automatic bootstrapping and configuration belong to this scope.

It is proposed that these solution scopes be addressed primarily sequentially by a CDNI WG and that the initial charter be centered around the "Base CDNI" scope. However there is obvious benefit from having a solution for the "Base CDNI" scope that is amenable to extension for support of the "Operational" scope and "Enhanced" scope. Therefore it is proposed that the initial CDNI WG charter also includes definition of (at least) the main requirements for the "Operationalized CDNI" scope and "Enhanced CDNI" Scope, so those can be kept in mind when defining the solution for the "Base CDNI" scope.

6. Gap Analysis of relevant Standardization and Research Activities

There are a number of other standards bodies and industry forums that are working in areas related to CDN, and in some cases related to CDNI. This section will first outline the key standardization organizations undertaking related work, some related research projects, and will then outline any potential overlap with the proposed CDNI WG and any component that could potentially be reused by CDNI .

6.1. Related standardization activities

6.1.1. 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 WG charter [CDI-Charter]:

"

- The goal of this working group is to define protocols to allow the interoperation of separately-administered content networks.
- 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.
- 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.
- In its initial form, the working group is not chartered to deliver those protocols [...]

"

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

The CDI WG published 3 Informational RFCs:

- RFC 3466 [RFC3466] - "A Model for Content Internetworking (CDI)".
- RFC 3568 [RFC3568] - "Known Content Network (CN) Request-Routing Mechanisms".
- RFC 3570 [RFC3570] - "Content Internetworking (CDI) Scenarios".

6.1.2. 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 is preparing an extended version for Release 10, which is scheduled to be published later 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 Interconnect environment. [3GP-DASH] could also be a candidate for content acquisition between CDNs in a CDN Interconnect environment.

6.1.3. 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 Interconnect 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 Interconnect use cases.

6.1.4. 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.

6.1.5. 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.

6.1.6. 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 Interconnect.

6.1.7. ETSI TISPAN

ETSI TISPAN has published two sets of IPTV specifications, one of which is based on IMS. In addition, TISPAN is about to complete the specifications of 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. In line with its long history of leveraging IETF protocols, ETSI could potentially leverage CDNI protocols developed in the IETF for their related protocol level work on interconnections of CDNs.

6.1.8. 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.

6.1.9. 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:

- Media Formats including HTTP Adaptive Streaming
- Content Metadata
- Protocols
- Terminal (Declarative or Procedural Application Environment)
- Authentication, Content Protection and Service Protection

6.1.10. 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.

6.1.11. 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."

6.2. Related Research Projects

6.2.1. 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].

6.2.2. 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.

6.2.3. 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 made publically available as a deliverable of this project.

6.3. Gap Analysis

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

- TISPAN has a dedicated specification for CDN.
- OIPF and ATIS 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.
- <TODO: Add a sentence on ITU>
- IETF CDN WG (now concluded) touched on the same problem space as the present document. However, in accordance with its initial charter, the CDI WG 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 WG 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 protocols and CDNI Metadata representations required to enable even a base level of CDN Interconnect functionality to be implemented.

The following sections will summarize the existing work described in Section 6.1 against the CDNI problem space.

6.3.1. 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 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, :

- The OIPF data plane includes both RTP and HTTP flavors (HTTP progressive download, HTTP Adaptive streaming [3GP-DASH],...).
- ATIS 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 Interconnect environment. Therefore for the purpose of a CDNI WG 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. OIPTF data plane mentioned above, HTTP adaptive streaming [3GP-DASH]) or public specifications (e.g. vendor specific HTTP Adaptive streaming specification) so that content delivery is 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 CDNI. CDNI may only concern itself with the negotiation/selection aspects of the acquisition protocol to be used in a CDN interonnect scenario.

6.3.2. CDNI Metadata

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

- 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-Metadata]. 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.
- 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.
- 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 need of a CDN Interconnect solution, it is worth considering CDMI as one of the existing pieces of work that may potentially be leveraged for the CDNI Metadata protocol (e.g by extending the CDMI metadata to address more specific CDNI needs).

7. Relationship to relevant IETF Working Groups

7.1. ALTO

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 WG 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). See [I-D.penno-alto-cdn] for a detailed discussion on how CDN Request Routing can 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 Interconnect. 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 work of ALTO is complementary to and does not overlap with the work proposed in this document because the integration between ALTO and a CDN would fall under "algorithms for selection of CDN or Surrogate by Request-Routing systems" in Section 3.2 and is therefore out of scope for a CDNI WG. One area for further study is whether additional information should be provided by an ALTO service to facilitate CDNI CDN selection.

7.2. DECADE

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:

- 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
- 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 protocols in the neighbor CDNs), we observe that:

- CDNI would operate as a "Content Distribution Application" from the DECADE viewpoint (i.e. would operate on top of DECADE).
- 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).
- There would typically be limited benefits in making use of a DECADE in-network storage service because the CDNI protocols 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 proposed in this document.

7.3. PPSP

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 WG 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 proposed in this document for CDNI.

8. IANA Considerations

This document makes no request of IANA.

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

9. 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 users in line with the CSP policy. 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 (i.e. the CSP "trusts" a CDN that "trusts" another CDN).

Maintaining the security of the content itself, its associated metadata (including distribution and delivery policies) and the CDNs distributing and delivering it, are critical requirements for both CDN Providers and CSPs and any work on CDN Interconnection 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 CDN Interconnects.

10. Acknowledgements

The authors would like to thank Andre Beck, Mark Carlson, Bruce Davie, David Ferguson, Yiu Lee, Kevin Ma, Julien Maisonneuve, Guy Meador, Emile Stephan, Oskar van Deventer and Mahesh Viveganandhan for their review comments and contributions to the text.

11. References

11.1. Normative References

[I-D.bertrand-cdni-experiments]
Bertrand, G., Faucheur, F., and L. Peterson, "Content Distribution Network Interconnection (CDNI) Experiments", draft-bertrand-cdni-experiments-00 (work in progress), February 2011.

[I-D.bertrand-cdni-use-cases]
Bertrand, G., Stephan, E., Watson, G., Burbridge, T., and P. Eardley, "Use Cases for Content Distribution Network Interconnection", draft-bertrand-cdni-use-cases-01 (work in progress), January 2011.

[I-D.bertrand-cdni-use-cases-00]
Bertrand, G. and E. Stephan, "Use Cases for Content Distribution Network Interconnection - draft-bertrand-cdni-use-cases-00 (superseded)",

January 2011.

[I-D.jenkins-cdni-names]
Niven-Jenkins, B., "Thoughts on Naming and Referencing of Data Objects within Content Distribution Network Interconnection (CDNI) solutions", draft-jenkins-cdni-names-00 (work in progress), February 2011.

[I-D.watson-cdni-use-cases]
Watson, G., "CDN Interconnect Use Cases", draft-watson-cdni-use-cases-00 (work in progress), January 2011.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

11.2. 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 ht tp://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.penno-alto-cdn]
Penno, R., Raghunath, S., Medved, J., Alimi, R., Yang, R., and S. Previdi, "ALTO and Content Delivery Networks", draft-penno-alto-cdn-02 (work in progress), October 2010.

[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.

[RFC3920] Saint-Andre, P., Ed., "Extensible Messaging and Presence

Protocol (XMPP): Core", RFC 3920, October 2004.

[RFC5023] Gregorio, J. and B. de hOra, "The Atom Publishing Protocol", RFC 5023, October 2007.

[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 Recomendation Y.1910 "IPTV functional architecture"", September 2008.

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

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

Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 10, 2012

K. Leung
Cisco
Y. Lee
Comcast
F. Le Faucheur
M. Viveganandhan
Cisco
G. Watson
BT

July 9, 2011

Content Distribution Network Interconnection (CDNI) Requirements
draft-lefaucheur-cdni-requirements-02

Abstract

Content Delivery Networks (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. There is a requirement for interconnecting standalone CDNs so that their collective CDN footprint can be leveraged for the end-to-end delivery of content from Content Service Providers (CSPs) to end users. The Content Distribution Network Interconnection (CDNI) working group has been chartered to develop an interoperable and scalable solution for such CDN interconnection.

The goal of the present document is to outline the requirements for the solution and interfaces to be specified by the CDNI working group.

Requirements Language

The key words "Must", "Should" and "May" in this document are to be interpreted in the following way:

- "Must" indicates requirements that are to be supported by the CDNI protocols in the stated scope (aka "within initial CDNI scope" or "beyond initial scope"). A requirement is stated as a "Must" when it is established by that it can be met without compromising the targeted schedule for WG deliverables, or when it is established that specifying a solution without meeting this requirement would not make sense and would justify re-adjusting the WG schedule, or both.

- "Should" indicates requirements that are to be supported by the CDNI protocols in the stated scope (aka "within initial CDNI scope" or "beyond initial scope") unless the WG realizes at a

later stage that attempting to meet this requirement would compromise the overall WG schedule (for example it would involve complexities that would result in significantly delaying the deliverables).

- o "May" indicates requirements that are to be supported by the CDNI protocols in the stated scope (aka "within initial CDNI scope" or "beyond initial scope") provided that dedicating WG resources to this work does not prevent addressing "Should" and "Must" requirements and that attempting to meet this requirement would not compromise the overall WG schedule.

Status of this Memo

Copyright Notice

Table of Contents

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, 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 End User regardless of that End User's location or attachment network. However, the footprint of a given CDN in charge of delivering a given content may not expand close enough to the End User's current location or attachment network to realize the cost benefit and user experience that a more distributed CDN would provide. This creates a requirement for interconnecting standalone CDNs so that their collective CDN footprint can be leveraged 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.

[I-D.jenkins-cdni-problem-statement] outlines the problem area that the CDNI working group is chartered to address. [I-D.bertrand-cdni-use-cases] discusses the use cases for CDN Interconnection. [I-D.davie-cdni-framework] discusses the technology framework for the CDNI solution and interfaces.

The goal of the present document is to document the requirements for the CDNI solution and interfaces. In accordance with the working group charter, the work is prioritized in a "walk before you run" approach: the present document separates the CDNI requirements into a set of more urgent requirements that are within the initial scope of the CDNI working group, and a set of less urgent additional requirements that are left to potential future rechartering of the working group.

1.1. Terminology

This document uses the terminology defined in section 1.1 of [I-D.jenkins-cdni-problem-statement].

This also defined the following additional terms [Editor's Note: these definitions may be better located in another document such as the Problem Statement]:

- Recursive CDNI request routing: When an Upstream CDN elects to redirect a request towards a Downstream CDN, the Upstream CDN can query the Downstream CDN Request Routing system via the CDNI Request Routing protocol (or use information cached from earlier similar queries) to find out how the Downstream CDN wants the request to be redirected, which allows the Upstream CDN to factor in the Downstream CDN response when redirecting the user agent. This approach is referred to as "recursive" CDNI request routing. Note that the Downstream CDN may elect to have the request redirected directly to a Surrogate inside the Downstream CDN, to the Request-Routing System of the Downstream CDN, to another CDN, or to any other system that the Downstream CDN sees as fit for handling the redirected request.

- Iterative CDNI Request Routing: When an Upstream CDN elects to redirect a request towards a Downstream CDN, the Upstream CDN can base its redirection purely on a local decision (and without attempting to take into account how the Downstream CDN may in turn redirect the user agent). In that case, the Upstream CDN redirects the request to the request routing system in the Downstream CDN, which in turn will decide how to redirect that request: this approach is referred to as "iterative" CDNI request routing.

2. CDNI Model and CDNI protocols

For convenience Figure 1 from [I-D.jenkins-cdni-problem-statement] illustrating the CDNI problem area and the CDNI protocols is replicated below.

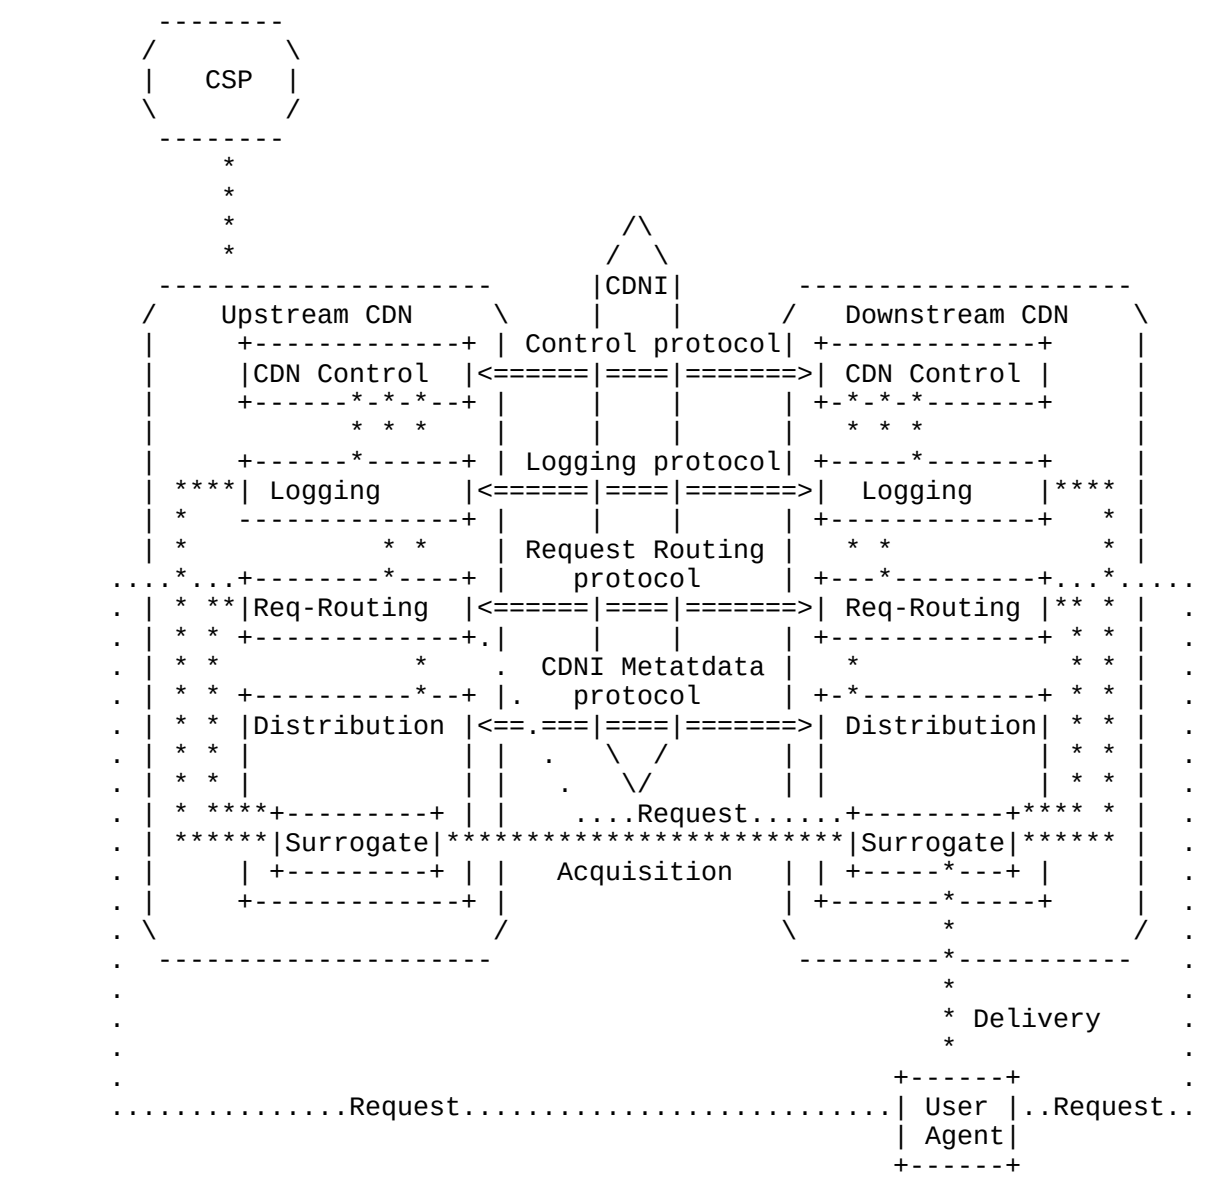

<==> interfaces inside the scope of CDNI

**** interfaces outside the scope of CDNI
.... interfaces outside the scope of CDNI

Figure 1: CDNI Model and CDNI APIs

3. Generic Requirements

This section identifies generic requirements independent of the individual CDNI protocols. Some of those are expected to affect multiple or all protocols.

3.1. Within Initial CDNI Scope

R1 Wherever possible, the CDNI protocols Should reuse or leverage existing IETF protocols.

R2 The CDNI solution Must not require a change, or an upgrade, to the User Agent to benefit from content delivery through interconnected CDNs.

R3 The CDNI solution Must not require intra-CDN information to be exposed to other CDNs for effective and efficient delivery of the content. Examples of intra-CDN information include surrogate topology, surrogate status, cached content, etc.

R4 The CDNI solution Must support delivery to the user agent based on HTTP [RFC2616]. [Note that while delivery and acquisition "data plane" protocols are out of the CDNI solution scope, the CDNI solution "control plane" protocols are expected to participate in enabling, selecting or facilitating operations of such acquisition and delivery protocols. Hence it is useful to state requirements on the CDNI solution in terms of which acquisition and delivery protocols].

R5 The CDNI solution Must support acquisition across CDNs based on HTTP [RFC2616].

R6 The CDNI solution May support delivery to the user agent based on protocols other than HTTP.

R7 The CDNI solution May support acquisition across CDNs based on protocols other than HTTP.

R8 The CDNI solution Should support cascaded CDN redirection (CDN1 redirects to CDN2 that redirects to CDN3) to an arbitrary number of levels.

R9 The CDNI solution Should support an arbitrary topology of interconnected CDNs (i.e. the CDN topology cannot be restricted to a tree, a loop-free topology, etc.).

R10 The CDNI solution Must prevent looping of any CDNI information exchange.

R11 When making use of third party reference, the CDNI solution Must consider the potential issues associated with the use of various format of third-party references (e.g. NAT or IPv4/IPv6 translation potentially breaking third-party references based on an IP addresses such as URI containing IPv4 or IPv6 address litterals, split DNS situations potentially breaking third-party references based on DNS fully qualified domain names) and wherever possible avoid, minimize or mitigate the associated risks based on the specifics of the environments where the reference is used (e.g. likely or unlikely presence of NAT in the path). In particular, this applies to situations where the CDNI solution needs to construct and convey uniform resource identifiers for directing/redirecting a content request, as well as to situations where the CDNI solution needs to pass on a third party reference (e.g. to identify a User Agent) in order to allow another entity to make a more informed decision (e.g. make a more informed request routing decision by attempting to derive location information from the third party reference).

3.2. Beyond Initial CDNI Scope

R12 The CDNI solution Must support cascaded CDN redirection (CDN1 redirects to CDN2 that redirects to CDN3) to an arbitrary number of levels. [Note: this "Must" requirement appeared as a "Should" requirement in Section 3.1]

R13 The CDNI solution Must support an arbitrary topology of interconnected CDNs (i.e. the CDN topology cannot be restricted to a tree, a loop-free topology, etc.). [Note: this "Must" requirement appeared as a "Should" requirement in Section 3.1]

R14 The CDNI solution Should support virtualization of the Downstream CDN, so that the Downstream CDN can appear as multiple logical Downstream CDNs.

4. CDNI Control Protocol Requirements

The primary purpose of the CDNI Control protocol is to initiate the interconnection across CDNs, bootstrap the other CDNI interfaces and trigger actions into the Downstream CDN by the Upstream CDN (such as delete object from caches or trigger pre-positioned content acquisition). We observe that while the CDNI Control protocol is currently discussed as a single "protocol", further analysis will determine whether the corresponding requirements are to be realized

over a single interface and protocol, or over multiple interfaces and protocols.

4.1. Within Initial CDNI Scope

R15 The CDNI Control protocol Must allow the Upstream CDN to request that the Downstream CDN (and, if cascaded CDNs are supported by the solution, that the potential cascaded Downstream CDNs) perform the following actions on an object or object set:

* Mark an object(s) and/or its CDNI metadata as "stale" and revalidate them before they are delivered again

* Delete an object(s) and/or its CDNI metadata from the CDN surrogates and any storage.

R16 The CDNI Control protocol Must allow the downstream CDN to report on the completion of these actions (by itself, and if cascaded CDNs are supported by the solution, by potential cascaded Downstream CDNs), in a manner appropriate for the action (e.g. synchronously or asynchronously).

R17 The CDNI Control protocol Must support initiation and control by the Upstream CDN of pre-positioned CDNI metadata acquisition by the Downstream CDN.

R18 The CDNI Control protocol Should support initiation and control by the Upstream CDN of pre-positioned content acquisition by the Downstream CDN.[Editor’s Note: how much influence the Upstream CDN ought to have on pre-positioning of the content on surrogates inside the Downstream CDN is TBD].

4.2. Beyond Initial CDNI Scope

R19 The CDNI Control protocol Must support support initiation and control by the Upstream CDN of pre-positioned content acquisition.[Editor’s Note: how much influence the Upstream CDN ought to have on pre-positioning of the content on surrogates inside the Downstream CDN is TBD]. [Note: this "Must" requirement appeared as a "Should" requirement in Section 4.1]

R20 The CDNI Control protocol Must allow a CDN to establish, update and terminate a CDN interconnection with another CDN whereby one CDN can act as a Downstream CDN for the other CDN (that acts as an Upstream CDN).

R21 The CDNI Control protocol Must allow control of the CDNI interconnection between any two CDNs independently for each direction (i.e. For the direction where CDN1 is the Upstream CDN and CDN2 is the Downstream CDN, and for the direction where CDN2 is the Upstream CDN and CDN1 is the Downstream CDN).

R22 The CDNI Control protocol Should allow bootstrapping of the Request-Routing protocol. For example, this can potentially include:

* negotiation of the Request-Routing method (e.g. DNS vs HTTP, if more than one method is specified)

* discovery of the Request-Routing protocol endpoints

* information necessary to establish secure communication between the Request-Routing protocol endpoints.

R23 The CDNI Control protocol Should allow bootstrapping of the Metadata Signaling protocol. This information could, for example, include:

* discovery of the Metadata Signaling protocol endpoints

* information necessary to establish secure communication between the Metadata Signaling protocol endpoints.

R24 The CDNI Control protocol Should allow bootstrapping of the Content Acquisition protocol. This could, for example, include exchange and negotiation of the Content Acquisition protocols to be used across the CDNs (e.g. HTTP, HTTPS, FTP, ATIS C2).

R25 The CDNI Control protocol Should allow exchange and negotiation of delivery authorization mechanisms to be supported across the CDNs (e.g. URI signature based validation).

R26 The CDNI Control protocol Should allow bootstrapping of the CDNI Logging protocol. This information could, for example, include:

* discovery of the Logging protocol endpoints

* information necessary to establish secure communication between the Logging protocol endpoints

* negotiation/definition of the log file format and set of fields to be exported through the Logging protocol, with some granularity (e.g. On a per content type basis).

* negotiation/definition of parameters related to transaction Logs export (e.g., export protocol, file compression, export frequency, directory).

5. CDNI Request Routing Protocol Requirements

5.1. Within Initial CDNI Scope

The main function of the Request Routing protocol is to allow the Request-Routing systems in interconnected CDNs to communicate to facilitate redirection of the request across CDNs.

R27 The CDNI Control protocol Must allow the Downstream CDN to communicate to the Upstream CDN coarse information about the Downstream CDN ability and/or willingness to handle requests from the Upstream CDN. For example, this could potentially include a binary signal ("Downstream CDN ready/not-ready to take additional requests from Upstream CDN") to be used in case of excessive load or failure condition in the Downstream CDN.

R28 The CDNI Request-Routing protocol Should allow the Downstream CDN to communicate to the Upstream CDN aggregate information to facilitate CDN selection during request routing, such as Downstream CDN capabilities, resources and affinities (i.e. Preferences or cost). This information could, for example, include:

* supported content types and delivery protocols

* footprint (e.g. layer-3 coverage)

* a set of metrics/attributes (e.g. Streaming bandwidth, storage resources, distribution and delivery priority)

* a set of affinities (e.g. Preferences, indication of distribution/delivery fees)

* information to facilitate request redirection (e.g. Reachability information of Downstream CDN Request Routing system).

[Note: Some of this information - such as supported content types and delivery protocols- may also potentially be taken into account by the distribution system in the Upstream CDN for pre-positioning of content and/or metadata in the Downstream CDN in case of pre-positioned content acquisition and/or pre-positioned CDNI metadata acquisition.]

R29 If cascaded redirection is supported by the CDNI solution, the CDNI Request-Routing protocol Must allow the Downstream CDN to also include in the information communicated to the Upstream CDN, information on the capabilities, resources and affinities of CDNs to which the Downstream CDN may (in turn) redirect requests received by the Upstream CDN. In that case, the CDNI Request-Routing protocol Must prevent looping of such information exchange.

R30 The CDNI Control protocol May allow the Downstream CDN to communicate to the Upstream CDN aggregate information on CDNI administrative limits and policy. This information can be taken into account by the Upstream CDN Request Routing system in its CDN Selection decisions. This information could, for example, include:

* maximum number of requests redirected by the Upstream CDN to be served simultaneously by the Downstream CDN

* maximum aggregate volume of content (e.g. in Terabytes) to be delivered by the Downstream CDN over a time period.

R31 The CDNI Request-Routing architecture and protocol Must support efficient request-routing for small objects. This may, for example, call for a mode of operation (e.g. DNS-based request routing) where freshness and accuracy of CDN/Surrogate selection can be traded-off against reduced request-routing load (e.g. Via lighter-weight queries and caching of request-routing decisions).

R32 The CDNI Request-Routing architecture and protocol Must support efficient request-routing for large objects. This may, for example, call for a mode of operation (e.g. HTTP-based request routing) where freshness and accuracy of CDN/Surrogate selection justifies a per-request decision and a per-request CDNI Request-Routing protocol call.

R33 The CDNI Request-Routing architecture Must support recursive CDNI request routing.

R34 The CDNI Request-Routing architecture Must support iterative CDNI request routing.

R35 In case of detection of a request redirection loop, the CDNI Request-Routing loop prevention mechanism Should allow routing of the request (as opposed to the request loop being simply interrupted without routing the request).

R36 The CDNI Request-Routing protocol Should support an optional mechanism allowing enforcment of a limit on the number of successive CDN redirections for a given request.

R37 The CDNI Request-Routing protocol May support an optional mechanism allowing an upstream CDN to avoid redirecting a request to a downstream CDN if that is likely to result in the total redirection time exceeding some limit.

R38 The CDNI Request-Routing protocol Must allow the Upstream CDN to include, in the query to the Downstream CDN, the necessary information to allow the Downstream CDN to process the redirection query. This could, for example, include:

* information from which the location of the user-agent that originated the request can be inferred (e.g. User Agent fully qualified domain name in case of HTTP-based Request Routing, DNS Proxy fully qualified domain name in case of DNS-based Request Routing)

* requested resource information (e.g. Resource URI in case of HTTP-based Request Routing, Resource hostname in case of DNS-based Request Routing)

* additional available request information (e.g. request headers in case of HTTP-based Request Routing).

R39 The CDNI Request-Routing protocol May also allow the Upstream CDN to convey information pointing to CDNI metadata applicable (individually or through inheritance) to the requested content. For illustration, the CDNI metadata pointed to could potentially include metadata that is applicable to any content, metadata that is applicable to a content collection (to which the requested content belongs) and/or metadata that is applicable individually to the requested content.

R40 The CDNI Request-Routing protocol Must allow the Downstream CDN to include the following information in the response to the Upstream CDN:

* status code, in particular indicating acceptance or rejection of request (e.g. Because the Downstream CDN is unwilling or unable to serve the request). In case of rejection, an error code is also to be provided, which allows the Upstream CDN to react appropriately (e.g. Select another Downstream CDN, or serve the request itself)

* redirection information (e.g. Resource URI in case of HTTP-based Request Routing, equivalent of a DNS record in case of DNS-based Request Routing).

5.2. Beyond Initial CDNI Scope

R41 The CDNI Request-Routing protocol Must allow the Downstream CDN to communicate to the Upstream CDN aggregate information to facilitate CDN selection during request routing, such as Downstream CDN capabilities, resources and affinities (i.e. Preferences or cost). This information could, for example, include:

* supported content types and delivery protocols

* footprint (e.g. layer-3 coverage)

* a set of metrics/attributes (e.g. Streaming bandwidth, storage resources, distribution and delivery priority)

* a set of affinities (e.g. Preferences, indication of distribution/delivery fees)

* information to facilitate request redirection (e.g. Reachability information of Downstream CDN Request Routing system).

[Note: this "Must" requirement appeared as a "Should" requirement in Section 5.1]

R42 The CDNI Request-Routing protocol Must allow the Downstream CDN to also include in the information communicated to the Upstream CDN, information on the capabilities, resources and affinities of CDNs to which the Downstream CDN may (in turn) redirect requests received by the Upstream CDN. The CDNI Control protocol Must prevent looping of such information exchange. [Note: this "Must" requirement appeared as a conditional "Must" requirement in Section 5.1]

R43 The CDNI Request-Routing protocol Should allow the Downstream CDN to communicate to the Upstream CDN aggregate information on CDNI administrative limits and policy. This information can be taken into account by the Upstream CDN Request Routing system in its CDN Selection decisions. This information could, for example, include:

* maximum number of requests redirected by the Upstream CDN that to be served simultaneously by the Downstream CDN

* maximum aggregate volume of content (e.g. in Terabytes) to be delivered by the Downstream CDN over a time period

[Note: this "Should" requirement appeared as a "May" requirement in Section 5.1]

R44 The CDNI Request-Routing loop prevention mechanism Must allow routing of the request (as opposed to the request loop being simply interrupted without routing the request). [Note: this "Must" requirement appeared as a "Should" requirement in Section 5.1]

R45 The CDNI Request-Routing protocol Must support optional enforcement of a limit on the number of successive CDN redirections for a given request. [Note: this "Must" requirement appeared as a "Should" requirement in Section 5.1]

6. CDNI Metadata Distribution Protocol Requirements

The primary function of the CDNI Metadata Distribution protocol is to allow the Distribution system in interconnected CDNs to communicate to ensure Content Distribution Metadata with inter-CDN scope can be exchanged across CDNs. We observe that while the CDNI Metadata Distribution protocol is currently discussed as a single "protocol", further analysis will determine whether the corresponding requirements are to be realized over a single interface and protocol, or over multiple interfaces and protocols. For example, a subset of the CDNI metadata might be conveyed in-band along with the actual content acquisition across CDNs (e.g. content MD5 in HTTP header) while another subset might require an out-of-band interface & protocol (e.g. geo-blocking information).

6.1. Within Initial CDNI Scope

R46 The CDNI Metadata Distribution protocol Must allow the Upstream CDN to provide the Downstream CDN with content distribution metadata of inter-CDN scope.

R47 The CDNI Metadata Distribution protocol Must support exchange of CDNI metadata for both the dynamic content acquisition model and the pre-positioning content acquisition model.

R48 The CDNI Metadata Distribution protocol Must/Should/May? support a mode where no, or a subset of, the Metadata is initially communicated to the Downstream CDN along with information about how/where to acquire the rest of the CDNI Metadata (i.e. Dynamic CDNI metadata acquisition).

R49 The CDNI Metadata Distribution protocol Must/Should/May? support a mode where all the relevant Metadata is initially communicated to the Downstream CDN (i.e. Pre-positioned CDNI metadata acquisition).

R50 Whether in the pre-positioned content acquisition model or in the dynamic content acquisition model, the CDNI Metadata Distribution protocol Must provide the necessary information to allow the Downstream CDN to acquire the content from an upstream source (e.g. Acquisition protocol and Uniform Resource Identifier in Upstream CDN- or rules to construct this URI).

R51 The CDNI metadata Must allow signaling of one or more upstream sources, where each upstream source can be in the Upstream CDN, in another CDN, the CSP origin server or any arbitrary source designated by the Upstream CDN. Note that some upstream sources (e.g. the content origin server) may or may not be willing to serve the content to the Downstream CDN, if this policy is known to the upstream CDN then it may omit those sources when exchanging CDNI metadata.

R52 The CDNI Metadata Distribution protocol Must allow the Upstream CDN to request addition and modification of CDNI Metadata into the Downstream CDN.

R53 The CDNI Metadata Distribution protocol Must allow removal of obsolete CDNI Metadata from the Downstream CDN (this could, for example, be achieved via an explicit removal request from the Upstream CDN or via expiration of a Time-To-Live associated to the Metadata).

R54 The CDNI Metadata Distribution protocol Must allow association of CDNI Metadata at the granularity of individual object. This is necessary to achieve fine-grain Metadata distribution at the level of an individual object when necessary.

R55 The CDNI Metadata Distribution protocol Must allow association of CDNI Metadata at the granularity of an object set. This is necessary to achieve scalable distribution of metadata when a large number of objects share the same distribution policy.

R56 The CDNI Metadata Distribution protocol Must support multiple levels of inheritance with precedence to more specific metadata. For example, the CDNI Metadata Distribution protocol may support metadata that is applicable to any content, metadata that is applicable to a content collection and metadata that is applicable to an individual content where content level metadata overrides content collection metadata that overrides metadata

for any content.

R57 The CDNI Metadata Distribution protocol Must ensure that conflicting metadata with overlapping scope are prevented or deterministically handled.

R58 The CDNI Metadata Distribution protocol Must provide indication by the Downstream CDN to the Upstream CDN of whether the CDNI metadata (and corresponding future request redirections) is accepted or rejected. When rejected, the CDNI Metadata Distribution protocol Must allow the Downstream CDN to provide information about the cause of the rejection.

R59 The CDNI Metadata Distribution protocol Must allow signaling of content distribution control policies. For example, this could potentially 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)

* delegation whitelist/blacklist (i.e. Information defining which downstream CDNs the content may/may not be delivered through)

R60 The CDNI Metadata Distribution protocol Must allow signaling of authorization checks and validation that are to be performed by the surrogate before delivery. For example, this could potentially include:

* need to validate URI signed information (e.g. Expiry time, Client IP address).

6.2. Beyond Initial CDNI Scope

R61 The CDNI Metadata Distribution protocol Must support a mode where no, or a subset of, the Metadata is initially communicated to the Downstream CDN along with information about how/where to acquire the rest of the CDNI Metadata (i.e. Dynamic CDNI metadata acquisition). [Note: this "Must" requirement appeared as a "Must/Should/May?" requirement in Section 6.1]

R62 The CDNI Metadata Distribution protocol Must support a mode where all the relevant Metadata is initially communicated to the Downstream CDN (i.e.Pre-positioned CDNI metadata acquisition). [Note: this "Must" requirement appeared as a "Must/Should/May?" requirement in Section 6.1]

R63 The CDNI Metadata Distribution protocol Must allow signaling of CDNI-relevant surrogate cache behavior parameters. For example, this could potentially include:

* control of whether the query string of HTTP URI is to be ignored by surrogate cache

* content revalidation parameters (e.g. TTL)

7. CDNI Logging Protocol Requirements

This section identifies the requirements related to the CDNI Logging protocol. We observe that while the CDNI Logging protocol is currently discussed as a single "protocol", further analysis will determine whether the corresponding requirements are to be realized over a single interface and protocol, or over multiple interfaces and protocols.

7.1. Within Initial CDNI Scope

R64 The CDNI logging architecture and protocol Must ensure reliable logging of CDNI events.

R65 The CDNI Logging protocol Must provide logging of deliveries to User Agents performed by the Downstream CDN as a result of request redirection by the Upstream CDN.

R66 If cascaded CDNs are supported, the CDNI logging protocol Must allow the Downstream CDN to report to the Upstream CDN logging for deliveries performed by the Downstream CDN itself as well as logging for deliveries performed by cascaded CDNs on behalf of the Downstream CDN.

R67 The CDNI Logging protocol Must provide logging of distribution performed by the Upstream CDN as a result of acquisition request by the Downstream CDN.

R68 The CDNI Logging protocol Must support batch/offline exchange of logging records.

R69 The CDNI Logging protocol Should also support additional timing constraints for some types of logging records (e.g. near-real time for monitoring and analytics applications)

R70 The CDNI Logging protocol Must define a log file format and a set of fields to be exported through the Logging protocol, with some granularity (e.g. On a per content type basis).

R71 The CDNI Logging protocol Must define a transport mechanisms to exchange CDNI Logging files.

[Editor's note: should we add a requirement for support of aggregate/ summarized logs (e.g. total bytes delivered for a content regardless of individual USer Agents to which it was delivered)]

7.2. Beyond Initial CDNI Scope

R72 The CDNI logging protocol Must allow the Downstream CDN to report to the Upstream CDN logging for deliveries performed by the Downstream CDN itself as well as logging for deliveries performed by cascaded CDNs on behalf of the Downstream CDN. [Note: this "Must" requirement appeared as a conditional "Must" requirement in Section 7.1]

R73 The CDNI Logging protocol Must support real-time exchange of some types of logging records (e.g. For real-time monitoring of deliveries across CDNs). [Note: this "Must" requirement appeared as a "Should" requirement in Section 7.1]

R74 The CDNI Logging protocol Must allow a CDN to query another CDN for relevant current logging records (e.g. For on-demand access to real-time logging information).

8. CDNI Security Requirements

This section identifies the requirements related to the CDNI security. Some of those are expected to affect multiple or all protocols.

8.1. Within Initial CDNI Scope

R75 All the CDNI protocols Must support secure operation over unsecured IP connectivity (e.g. The Internet). This includes authentication, confidentiality, integrity protection as well as protection against spoofing and replay.

R76 The CDNI solution Must provide sufficient protection against Denial of Service attacks. This includes protection against spoofed delivery requests sent by user agents directly to a Downstream CDN attempting to appear as if they had been redirected by a given Upstream CDN when they have not.

R77 The CDNI solution Should be able to ensure that for any given request redirected to a Downstream CDN, the chain of CDN Delegation (leading to that request being served by that CDN) can be established with non-repudiation.

R78 The CDNI solution Should be able to ensure that the Downstream CDN cannot spoof a transaction log attempting to appear as if it corresponds to a request redirected by a given Upstream CDN when that request has not been redirected by this Upstream CDN. This ensures non-repudiation by the Upstream CDN of transaction logs generated by the Downstream CDN for deliveries performed by the Downstream CDN on behalf of the Upstream CDN.

R79 The CDNI solution May provide a mechanism allowing an Upstream CDN that has credentials to acquire content from the CSP origin server (or another CDN), to allow establishment of credentials authorizing the Downstream CDN to acquire the content from the CSP origin server (or the other CDN) (e.g. In case the content cannot be acquired from the Upstream CDN).

8.2. Beyond Initial CDNI Scope

R80 The CDNI solution Must provide a mechanism allowing an Upstream CDN that has credentials to acquire content from the CSP origin server (or another CDN), to allow establishment of credentials authorizing the Downstream CDN to acquire the content from the CSP origin server (or the other CDN) (e.g. In case the content cannot be acquired from the Upstream CDN). [Note: this "Must" requirement appeared as a "May" requirement in Section 8.1]

9. IANA Considerations

This document makes no request of IANA.

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

10. Security Considerations

This document discusses CDNI security requirements in Section 8.

11. Acknowledgements

This document leverages the earlier work of the IETF CDI working group in particular as documented in [I-D.cain-request-routing-req], [I-D.amini-cdi-distribution-reqs] and [I-D.gilletti-cdnp-aaa-reqs].

The authors would like to thank Gilles Bertrand, Christophe Caillet, Bruce Davie, Phil Eardly, Agustin Schapira and Emile Stephan for their input. We also want to thank Ben Niven-Jenkins for his review and comments.

12. References

12.1. Normative References

[I-D.bertrand-cdni-use-cases]
Bertrand, G., Stephan, E., Watson, G., Burbridge, T., Eardley, P., and K. Ma, "Use Cases for Content Delivery Network Interconnection", draft-bertrand-cdni-use-cases-02 (work in progress), July 2011.

[I-D.davie-cdni-framework]
Davie, B. and L. Peterson, "Framework for CDN Interconnection", draft-davie-cdni-framework-00 (work in progress), July 2011.

[I-D.jenkins-cdni-problem-statement]
Niven-Jenkins, B., Faucheur, F., and N. Bitar, "Content Distribution Network Interconnection (CDNI) Problem Statement", draft-jenkins-cdni-problem-statement-02 (work in progress), March 2011.

[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.

12.2. Informative References

[I-D.amini-cdi-distribution-reqs]
Amini, L., "Distribution Requirements for Content Internetworking", draft-amini-cdi-distribution-reqs-02 (work in progress), November 2001.

[I-D.cain-request-routing-req]
Cain, B., "Request Routing Requirements for Content Internetworking", draft-cain-request-routing-req-03 (work in progress), November 2001.

[I-D.gilletti-cdnp-aaa-reqs]
"CDI AAA Requirements,
draft-gilletti-cdnp-aaa-reqs-01.txt", June 2001.

Authors' Addresses

Kent Leung
Cisco Systems
3625 Cisco Way
San Jose 95134
USA

Phone: +1 408 526 5030
Email: kleung@cisco.com

Yiu Lee
Comcast

Email: yiu_lee@cable.comcast.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

Mahesh Viveganandhan
Cisco Systems
375 East Tasman Drive
San Jose 95134
USA

Email: mvittal@cisco.com

Grant Watson
BT

Email: grant.watson@bt.com

Network Working Group K. Ma
Internet-Draft R. Nair
Intended status: Informational Azuki Systems, Inc.
Expires: September 8, 2011 March 7, 2011

Content Distribution Network Interconnection (CDNI) Publisher Use Cases draft-ma-cdni-publisher-use-cases-00

Abstract

Content publishers are increasingly using multiple CDNs to publish content to the consumers. It is important to take into account their specific requirements with respect to workflow restrictions with respect to content licensing rules. Also, certain content applications have specific delivery requirements that need to considered while negotiating or signaling inter-CDN arrangements. This contribution highlights some of these use cases to help motivate discussion.

Requirements Language

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 RFC 2119 [RFC2119].

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 July 30, 2011.

Copyright Notice

Table of Contents

1. Introduction

Many use cases for CDN Interconnection are discussed in [I-D.bertrand-cdni-use-cases] from the viewpoint of the CDSP. However, with the growing popularity of over-the-top (OTT) delivery, many content providers (CP) are now choosing to run their own service, without operating their own CDN. The CP may contract and manage multiple CDNs, as opposed to selecting an "Authoritative CDN" to manage CDN delegation, as described in [I-D.bertrand-cdni-use-cases]. This draft takes the viewpoint of these CPs using a federation of CDNs.

This document provides additional use cases addressing the needs of workflow management agents (WMA), as they manage multiple Authoritative CDNs. More detailed use cases for unique content delivery scenarios for CPs and end users (EU) are also provided.

1.1. Terminology

This document uses the terminology defined in section 1.1 of [I-D.jenkins-cdni-problem-statement] and [I-D.bertrand-cdni-use-cases].

2. CDNI Model and CDNI APIs

For convenience Figure 1 from [I-D.jenkins-cdni-problem-statement] illustrating the CDNI problem area and the CDNI APIs is replicated below with the addition of the workflow management component.

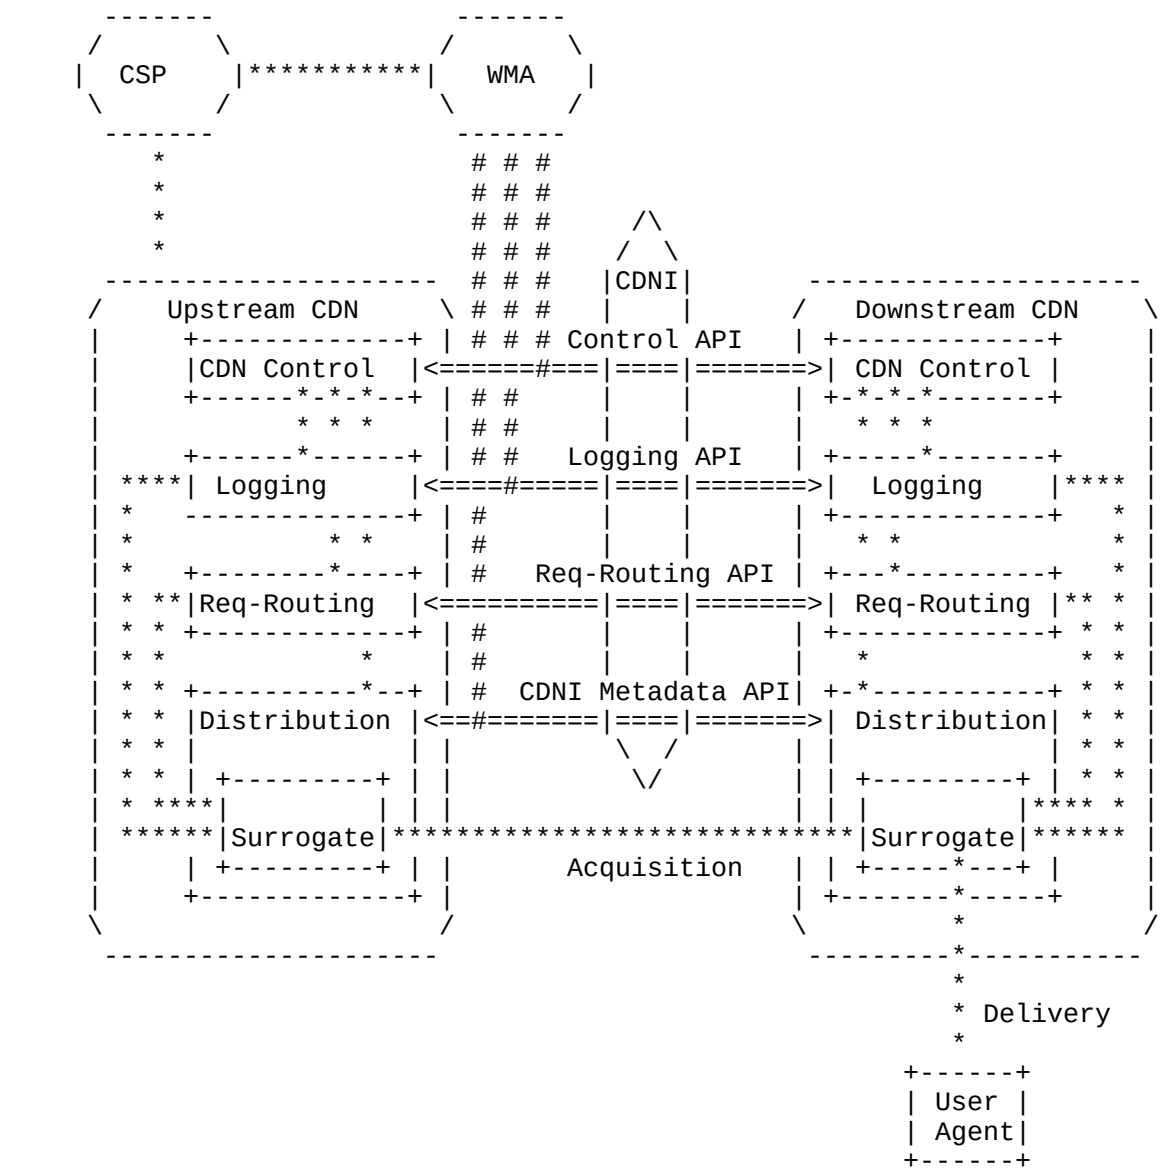

<==> interfaces inside the scope of CDNI

**** interfaces outside the scope of CDNI

interfaces important to Workflow Management Agents

Figure 1: CDNI Model and CDNI APIs

3. Workflow Management Use Cases

This section identifies generic requirements independent of the individual CDNI APIs. Some of those are expected to affect multiple or all APIs.

Content licensing is typically a complex mixture of temporal restrictions (e.g., how soon after the original release date), geo-location restrictions (e.g., in which countries), device and quality restrictions (e.g., only on devices with resolutions less than WxH), and NSP restrictions.

3.1. Asymmetric Content Availability Use Case

Often content is tailored to the NSP and devices supported by the NSP, e.g., NSP1 only support smart phones but not tablet devices, therefore only requires content with resolution less than 480x320. In this case, the content files available through the CDN (CDN1), operated by the NSP (NSP1), will be a subset of all available content files. A second NSP (NSP2) may support laptop and tablet devices and may therefore require higher resolution video. This CDN (CDN2), operated by NSP2, will require access to and possibly pre-positioning of different content files than CDN1. The ability to delegate requests from CDN2 to CDN1 is diminished by the availability of only a limited subset of the content files on CDN1. In this case, both CDN1 and CDN2 may be considered Authoritative CDNs, however, any resiliency measures introduced by CDN interconnection will be limited. The WMA must track and manage the rights of each Authoritative CDN to different content files and be able to set delegation policies for those CDNs.

3.2. Content Purge Use Case

Temporal licensing restrictions typically consist of a "sunrise" (activation time), a "sunset" (deactivation time), and a "takedown" (expiration time). Sunrise and sunset refer to logical availability, however, takedown refers to physical availability. Service agreements must typically support takedown guarantees. The ability to both issue and confirm the expunging of content files from surrogates is required to enforce these content licenses. These licenses may have different sunrise, sunset, and takedown times for different NSPs. The WMA must track, manage, and confirm the takedown of content from each of the Authoritative CDNs that content was

provided to. In the case that an Authoritative CDN delegated delivery to a non-Authoritative CDN, confirmation of content takedown from the non-Authoritative Downstream CDNs must be confirmed as well.

3.3. Service Level Agreement Negotiation Use Case

For video streaming services, the content provided for distribution is typically encoded at known bitrates. The quality of the delivered content is dependent upon the network capacity of the CDN. EU experience is critical to protecting CP brand loyalty. The content provider may wish to restrict delivery when insufficient bandwidth exists to deliver a high quality content viewing experience. The "Overload Handling" use case, discussed in [I-D.bertrand-cdni-use-cases], discusses expansion of bandwidth, but not the specific enforcement of minimum bandwidth or minimum latency requirements. Metadata needs to support the specification of these requirements. The WMA needs to be able to negotiate and set limitations for content delivery.

4. Detailed Content Delivery Use Cases

This section identifies generic requirements independent of the individual CDNI APIs. Some of those are expected to affect multiple or all APIs.

4.1. Grouped File Delegation

Delegating or pre-positioning content which are commonly requested in close temporal proximity is an important use case for CDN optimization and affects metadata concepts in CDN interconnection.

For the delivery of streaming video, the content may be separated into multiple files, e.g., segment files for HTTP Live Streaming (HLS) [I-D.pantos-http-live-streaming]. When delegating requests for HLS content, it may not be efficient to delegate each segment file request, but rather be able to delegate sets of content files that are commonly requested as a group. Web page images and icons may also fall into this category. For streaming video, however, there is also a sequential time component for which metadata may be used to optimize paced on-demand pre-fetching of segment files. Similarly, for the case of pre-positioning grouped files, the

ability to set a priority for certain files (e.g., common video starting points or hot-spots) allow the CDN to optimize storage usage, while still pre-positioning content critical to minimizing initial delivery latency. This information, combined with geo-location information, may enable the CDN to further optimize surrogate selection.

4.2. Adaptive Bitrate Streaming Use Case

Support for adaptive bitrate delivery an important use case for HTTP-based streaming video (e.g., HTTP Live Streaming (HLS) [I-D.pantos-http-live-streaming])

The CDN interconnection problem statement [I-D.jenkins-cdni-problem-statement] clearly states as a non-goal: "Content preparation, including encoding and transcoding". While the actual transformation of content is beyond the scope of CDN interconnection, the understanding about the existence of pre-transcoded files which are related to each other, is valuable for CDN pre-fetching, pre-positioning, and delivery optimization. This may be considered a further extension of the Grouped File Delegation use case. Metadata relaying the relative access patterns for specific content (e.g., switching from bitrate 6 to bitrate 7 is more likely than switching from bitrate 7 to bitrate 1) may be used to further optimize storage usage and surrogate selection.

4.3 Session Shifting Use Case

Session shifting between first (TV), second (PC), and third (mobile) screen devices is a primary use case for CDN interconnection.

The "Device and Network Technology Extension" use case, discussed in [I-D.bertrand-cdni-use-cases], covers a single device moving from one NSP/CDN to another (e.g., switching from WiFi to 3G on a mobile device). An extension to this scenario is switching between different devices either within the same NSP (e.g., switching from a TV on a cable network, to a PC connected to a cable modem or to a mobile device connected via WiFi to the cable modem) or across multiple NSPs (e.g., from a TV connected to a cable or a PC connected to a cable modem or a mobile device connected via WiFi to the cable modem, to a mobile device on a 3G/4G cellular connection). The requests may come from different devices and/or access networks, but all belong to the same virtual data stream.

4.4 Origin Server Failover Use Case

Support for multiple origin servers from which Downstream CDNs may retrieve content is an important use case for resilient delivery.

The "Overload Handling" and "Resiliency" use cases, discussed in [I-D.bertrand-cdni-use-cases], cover cases where an EU is redirected to an alternate CDN (CDN2) when the CDN which received the initial request (CDN1) is unable to service it due to capacity issues or other failure. An extension to this scenario is the condition where the origin server (OS1) becomes either temporarily or permanently unreachable to CDN1. In this case, there are two possible options:

- Redirect the request to CDN2, in the hope that CDN2 still has connectivity to the origin server (OS1).

- Failover to a secondary origin server (OS2) and attempt to retrieve the content from OS2 in order to service the request.

There may be a arbitrary number of alternate origin servers (which may in fact be other CDNs) and failover to an alternate CDN should only occur after exhausting a possibly limited alternate origin server search. In the case of live streaming video, an origin server may go down unexpectedly, and retrieval from an alternate CDN may not be feasible if it introduces too much latency.

5. Conclusions

This draft introduces some content publisher use cases that need to be considered in the CDNI efforts. We covered the cases that apply to workflow management as well as content delivery.

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

Content licensing rights need to be managed across multiple CDN domains as described in the workflow management use cases above.

7. Acknowledgements

The authors would like to thank the organizers of this effort to discuss the requirements for delivering premium content over the Internet.

8. References

8.1. Normative References

[I-D.bertrand-cdni-use-cases]
Bertrand, G. and E. Stephan, G. Watson, T. Burbridge, P. Eardley, "Use Cases for Content Distribution Network Interconnection", draft-bertrand-cdni-use-cases-01 (work in progress), January 2011.

[I-D.jenkins-cdni-problem-statement]
Niven-Jenkins, B., Faucheur, F., and N. Bitar, "Content Distribution Network Interconnection (CDNI) Problem Statement", draft-jenkins-cdni-problem-statement-01 (work in progress), January 2011.

[I-D.watson-cdni-use-cases]
Watson, G., "CDN Interconnect Use Cases", draft-watson-cdni-use-cases-00 (work in progress), January 2011.

[I-D.pantos-http-live-streaming]
R. Pantos, "HTTP Live Streaming", draft-pantos-http-live-streaming-05(work in progress), November 2011.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[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.

8.2. Informative References

[Apache-Common]
"Apache Common Log File Format (http://www.w3.org/Daemon/User/Config/Logging.html#common-logfile-format)".

[Apache-Format]
"Apache LogFormat Directive (http://httpd.apache.org/docs/1.3/mod/mod_log_config.html#logformat)".

[I-D.amini-cdi-distribution-reqs]
Amini, L., "Distribution Requirements for Content Internetworking", draft-amini-cdi-distribution-reqs-02 (work in progress), November 2001.

[I-D.cain-request-routing-req]
Cain, B., "Request Routing Requirements for Content Internetworking", draft-cain-request-routing-req-03 (work in progress), November 2001.

[I-D.gilletti-cdnp-aaa-reqs]
"CDI AAA Requirements, draft-gilletti-cdnp-aaa-reqs-01.txt", June 2001.

Authors' Addresses

Kevin Ma
Azuki Systems
43 Nagog Park
Acton, MA 01720
USA

Phone: +1 978-844-5100
Email: kevin.ma@azukisystems.com

Raj Nair
Azuki Systems
43 Nagog Park
Acton, MA 01720
USA

Phone: +1 978-844-5100
Email: raj.nair@azukisystems.com

Network Working Group
Internet-Draft
Intended status: Informational
Expires: September 30, 2011

B. Thompson
Cisco
A. Kobayashi
NEC

March 29, 2011

ATIS Internet Sourced Content Initiative and Relevance to CDNI
draft-thompson-cdni-atis-scenarios-00

Abstract

The ATIS IPTV Interoperability Forum (IIF) is a leading developer of requirements, standards, and specifications for Internet Protocol Television, or IPTV. IIF has specified an architecture for Content On Demand applications which was published as the [ATIS-0800042] IPTV Content on Demand (CoD) Service specification. ATIS IIF is now working on revision 2 of this document. Revision 2 of ATIS CoD specification includes an additional work item called "Internet Sourced Content" which includes "off net delivery". Both Internet Sourced Content and off net delivery will address scenarios that include delivery of content from a Content Service Provider through multiple CDNs operated by different organizations to a user agent. This document provides information on [ATIS-0800042], and the Internet Sourced Content and off net Delivery use cases that are relevant to CDNI.

The ATIS IIF architecture group focuses on developing architectures and prefers to reference existing protocols where ever possible. If, and as, the IETF CDNI effort progresses, the ATIS IIF architecture group will be evaluating the protocols defined by the CDNI working group for inclusion in the CDN interconnect scenarios defined as part of the Internet Sourced Content and off net Delivery use cases.

A goal of the present document is to illustrate the Internet Sourced Content and off net Delivery use cases defined in revision 2 of the ATIS CoD architecture that are relevant to the IETF CDNI work. It is hoped that these use cases can also be used to guide requirements for the development of IETF CDNI protocols.

Requirements Language

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 RFC 2119 [RFC2119].

Status of this Memo

Copyright Notice

Table of Contents

1. Introduction

[ATIS-0800042] specifies a CDN architecture in which the Origin Server (ATIS Content Origin Function), CDN Distribution System (ATIS Content Distribution and Delivery Functions), and User Agent (ATIS ITF) are all within the administrative domain of a Service Provider. In this respect, a Service Provider as defined in [ATIS-0800042] fulfills both the Content Service Provider and Network Service roles as defined in draft-jenkins-cdni-problem-statement-01.

The ATIS Internet Sourced Content Initative will build on the architecture specified in ATIS-0800042 to provides services across CDN domain boundaries. Many of the CDN functions and reference points specified in [ATIS-0800042] can be mapped to the achitecture specified in draft-jenkins-cdni-problem-statement-01. Section 2 below provides the mapping between the CDN architecture specified in [ATIS-0800042] and the architecture specified draft-jenkins-cdni-problem-statement-01. Sections 3 and 4 provide information about the off net Delivery and Internet Sourced Content use cases.

1.1. Terminology

This document uses the terminology defined in section 1.1 of [I-D.jenkins-cdni-problem-statement] and [ATIS-0800042].

2. ATIS CoD Mapping to CDNI

The sections below provide a mapping between the ATIS architecture and the architecture defined in [I-D.jenkins-cdni-problem-statement]

2.1. Domain Mapping

The ATIS IIF Internet Sourced Content service is based on the architecture specified in [ATIS-0800042] to provide services across provider domain boundaries. The provider domains defined as part of the ATIS Internet Sourced Content service include Content Provider, Service Provider, and Network Provider.

The ATIS Content Provider domain is similar to the CDNI Content Service Provider. One difference between the ATIS Content Provider and the CDNI Content Service Provider is that the ATIS Content Provider may only host the Origin Server for Content associated with services hosted by the ATIS Service Provider.

The ATIS Service Provider hosts services and also manages the access network, and provides a Content Service to End Users. In this

respect, the ATIS Service Provider may take on the roles of both the CDNI Content Service Provider and the Network Service Provider. An ATIS Service Provider typically manages their access network and uses tools such as QoS and policy management to ensure that content delivered over their access network is delivered with a guaranteed minimum level of quality.

The ATIS Network provider is the same as the CDNi Network provider.

The architectural diagrams in sections 3 and 4 illustrate ATIS use cases in terms of the roles defined in [I-D.jenkins-cdni-problem-statement]

2.2. Architectural Mapping

Figure 1 below illustrates the architecture and the reference points that are relevant to content distribution within [ATIS-08000042]. Figure 1 also illustrates the mapping of ATIS functions to the CDNI functions in [I-D.jenkins-cdni-problem-statement].

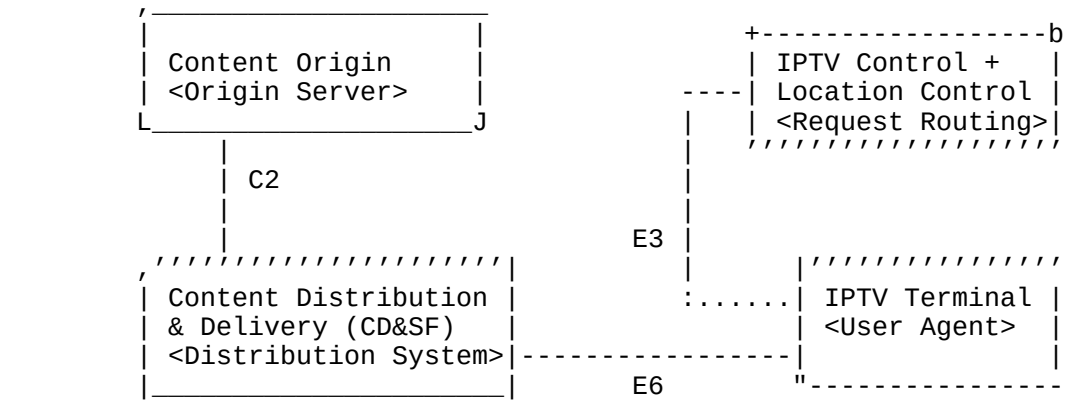

Figure 1: ATIS CDN Architecture

In Figure 1, the ATIS Content Origin Function is equivalent to an Origin Server as defined in RFC 2616. The ATIS Content Delivery and Storage Function (CD&SF) is equivalent to a CDNI Distribution System. The ATIS IPTV Terminal Function is equivalent to the CDNI User Agent function and the ATIS Location Control Function is equivalent to the CDNI Request Routing Function. The ATIS IPTV Control Function manages sessions initiated by IPTV Terminal function. The session interfaces maintained by the IPTV Control Function may be used to allocate bandwidth for continuous media delivered over the transport network using the services of a Resource and Admission Control (RACF)

function.

The ATIS IPTV Control + Location Control Functions are referenced from the IPTV Terminal Function via the E3 reference point. The ATIS E3 reference point provides the result of the request routing decision to the IPTV Terminal via either an HTTP or RTSP redirect.

The ATIS E6 reference point implements a stream control and content delivery interface. The interface definition of the E6 reference point specifies both http and rtsp variants.

The ATIS C2 reference point interfaces to the ATIS Content Origin Function and may also interface to a CDNI upstream CDN. When the ATIS C2 reference point is used to interface to a CDNI upstream CDN, it implements the CDNI request and aquisition interfaces. The ATIS C2 reference point defines 2 levels of functionality / compliance. They are the Basic Content Origin and the Extended Content Origin. The functionality of of these 2 levels of compliance is described below.

The Basic Content Origin Function of the ATIS C2 reference point defines metadata that may be considered to be inter-CDN Content Distribution Metadata in CDNI. An example ATIS C2 metadata that may be considered as inter-CDN Content Distribution Metadata is information about the bandwidth required to stream continuous media delivered as HTTP resources in real time. The ATIS Content Distribution & Delivery Function may use this information ensure that resources are reserved in order to ensure real time delivery to a CDNI downstream CDN or to User Agents using the services of the RACF function. Metadata delivered via the C2 reference point is referred to as Media Resource Metadata. The ATIS C2 reference point is based on http. When an http resource has Media Resource Metadata associated with it, the http link header provides a URL where the client can obtain the metadata. The interface definition for the ATIS C2 reference point includes an XML schema for Media Resource Metadata.

The Extended Content Origin Function of the ATIS C2 reference point defines HTTP extensions such as the Scheduled Transmission Service which is used to enable minimum bandwidth guarantees for continuous media delivered via HTTP. The ATIS C2 Scheduled Transmission Service enables an ATIS C2 client to request delivery of an HTTP resource at a specified minimum rate. An ATIS C2 server that supports the Scheduled Transmission Service may allocate resources associated with requests. An ATIS C2 server that does not have the resources to deliver a requested resource at the specified minimum rate will fail the requst due to lack of resources. Section 4.2 describes a use case where the use of the ATIS C2 Scheduled Transmission Service may

result in additional requirements to the CDNI protocols.

3. Off Net Delivery

In the Off Net use case, content hosted by an ATIS Service Provider is delivered to subscribers who are roaming. Content is delivered from the ATIS Service Provider's Origin Server and CDN through a CDN managed by a Network Service Provider to a user agent. In this use case, the ATIS Service Provider appears as both a CDNI Content Service Provider and a Network Service Provider which operates the CDNI upstream CDN, while the Network Service Provider operates the CDNI downstream CDN.

Figure 2 below illustrates the ATIS architecture and the ATIS reference points that are relevant to the Off Net Delivery use case. Figure 2 also illustrates the mapping of ATIS reference points to the CDNI interfaces in [I-D.jenkins-cdni-problem-statement].

In Figure 2, the User Agent originally requests content via the E3 reference point to the ATIS service provider CDN which appears as a CDNI upstream CDN. The E3 reference point implements the CDNI request interface. The ATIS Service provider CDN uses the CDNI request routing protocol to determine which downstream CDN to use for the requested content. The E3 reference point is then used to provide a URI which references the selected downstream CDN via an HTTP or RTSP redirect response.

This ATIS use case is expected to strictly match the CDNI problem statement defined in [I-D.jenkins-cdni-problem-statement] and its requirements are believed to be already covered in [I-D.lefaucheur-cdni-requirements].

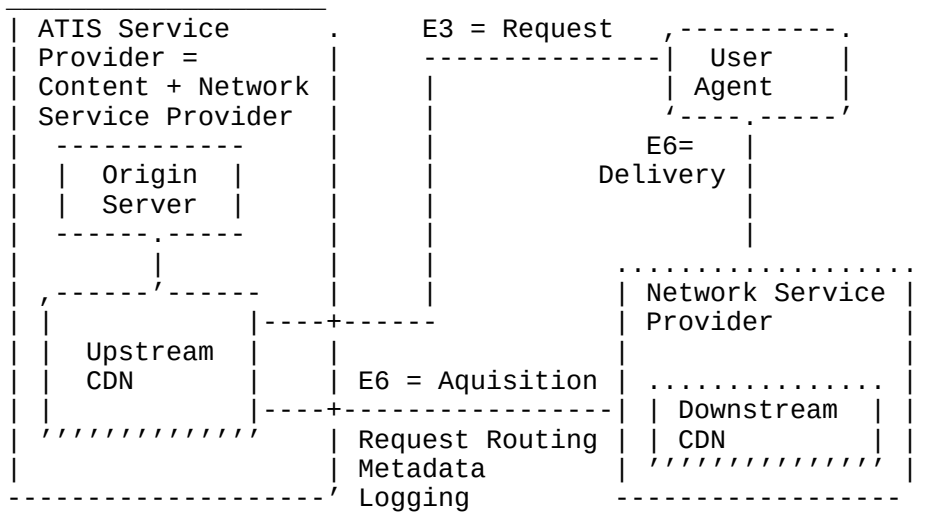

Figure 2: ATIS Off Net Delivery

4. Internet Sourced Content

In the Internet Sourced Content use case, content is delivered from an ATIS Content Provider to ATIS Service Provider subscribers over the ATIS Service Provider's CDN. Content may be delivered to the ATIS Service Provider's CDN directly from the ATIS Content Provider or through a CDN operated by a Network Service Provider that interconnects the ATIS Content Provider and the ATIS Service Provider CDN. In both cases, the Content is delivered from the ATIS Service Provider's CDN to the User Agent over a managed network. The transport network may include a RACF function. The RACF function may be used to allocate bandwidth for continuous media delivered over the transport network.

Two Internet Source Content use cases which use the CDNI APIs / Protocols are described below.

4.1. Transparent Operation

Figure 3 below illustrates the architecture and the reference points that are relevant to the Internet Sourced Content use case when the IETF CDNI APIs / protocols are transparent to the ATIS Service Provider's CDN. Figure 3 also illustrates the mapping of ATIS reference points to the CDNI interfaces in [I-D.jenkins-cdni-problem-statement].

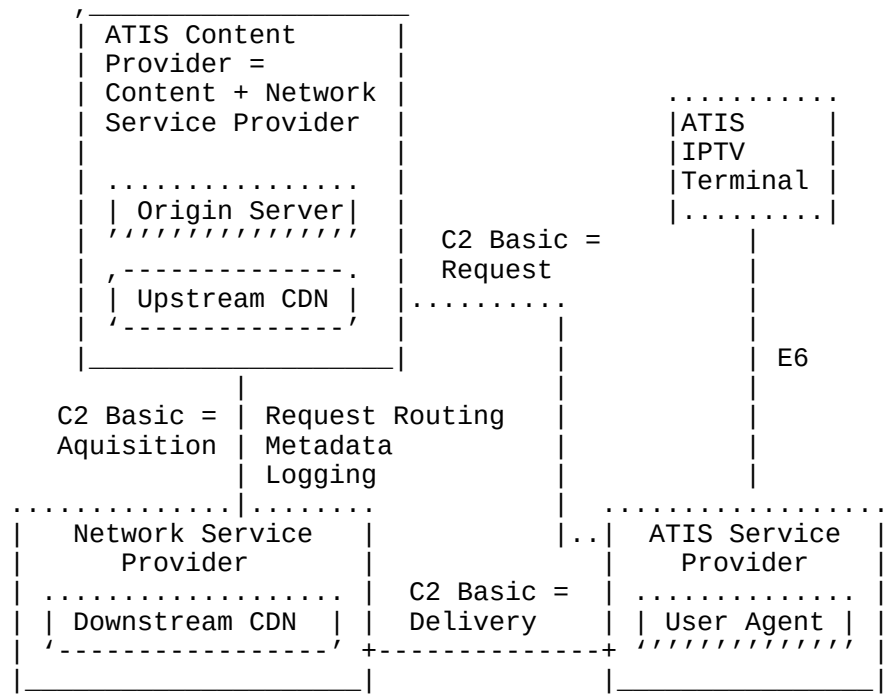

Figure 3: ATIS Internet Sourced Content Transparent Operation

In this use case, the ATIS Content Provider appears as both a CDNI Content Service Provider and a Network Service Provider which operates the CDNI upstream CDN. The Network Service Provider between the ATIS Content provider and ATIS Service Provider operates the CDNI downstream CDN. The ATIS Service Provider operates a CDN which appears as a User Agent to the Network Service Provider's Downstream CDN. The operation of the CDNI APIs / protocols is between the Content + Network Service Provider's Upstream CDN and the Network Service Provider's Downstream CDN. The ATIS Service Provider's CDN is therefore not using CDNI protocols and is not involved in the selection of the upstream CDN(s).

The ATIS Service Provider's CDN receives the initial resource request from the ATIS IPTV Terminal via the E6 reference point. On a cache miss, the ATIS Service Provider's CDN reflects the resource request to the Content Service Provider's Origin Server via the ATIS C2 reference point which acts as the CDNI Request Protocol. The Content Service Provider's CDN selects a downstream CDN for the request and the selected downstream CDN is provided to the ATIS Service Provider's CDN via the C2 reference point in a redirect response. When the ATIS Service Provider's CDN receives the redirect response,

it uses the returned URI to re-issues the ATIS C2 request to the selected Network Service Provider. This second ATIS C2 transaction therefore appears as the CDNI Delivery Protocol.

Since the delivery network between the ATIS Service Provider's CDN and the user agent may include a RACF function, the inherent bit rate of Continuous Media originating from the Content Service Provider should be provided for content delivered to the ATIS Service Provider's CDN. The currently accepted architecture for Internet Sourced Content uses the ATIS C2 reference point between the Content Service Provider, the Network Service Provider's upstream CDN, and the ATIS Service Provider's downstream CDN to deliver this information via Media Resource Metadata. Since Media Resource Metadata appears as an HTTP resource, the Network Service Provider's CDN may participate transparently in this architecture via regular HTTP caching.

This use case assumes there are enough resources from the Content Service Provider's Upstream CDN through the Network Service Provider's CDN to the ATIS Service Provider's CDN to enable the requested resource to be delivered at or above the rate required to deliver it to the ATIS IPTV Terminal in real time. The use case is representative of a scenario where the CDNs upstream of the ATIS Service Provider support either HTTP or the Basic Origin Functionality of ATIS C2.

This ATIS use case is expected to match the CDNI problem statement defined in [I-D.jenkins-cdni-problem-statement] and its requirements are believed to be already covered in [I-D.lefaucheur-cdni-requirements].

4.2. ATIS Service Provider Based CDN Selection

In this use case, the CDN path selected from the ATIS Content Provider to the ATIS Service Provider is based on criteria provided by the ATIS Service Provider.

Figure 4below illustrates the architecture and the reference points that are relevant to this use case. Figure 4 also illustrates the mapping of ATIS reference points to the CDNI interfaces in [I-D.jenkins-cdni-problem-statement].

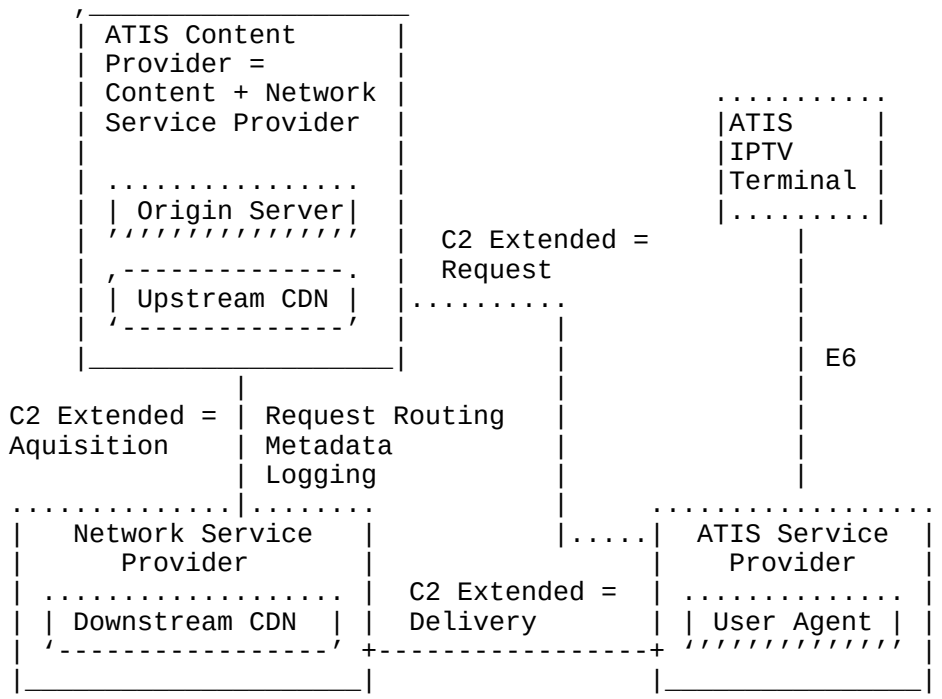

Figure 4: ATIS Service Provider Based CDN Selection

As in section 4.1, the ATIS Content Provider appears as both a CDNI Content Service Provider and a Network Service Provider which operates the CDNI upstream CDN. The Network Service Provider between the ATIS Content provider and ATIS Service Provider operates the CDNI downstream CDN. The ATIS Service Provider operates a CDN which appears as a User Agent to the Network Service Provider's Downstream CDN. The operation of the CDNI APIs / protocols is between the Content + Network Service Provider's Upstream CDN and the Network Service Provider's Downstream CDN.

The CDNI protocol interactions in this use case are almost identical to those described in section 4.1 except that the ATIS Service Provider uses the ATIS C2 reference point acting as the CDNI Request protocol to influence the CDNI path selection process. The ATIS C2 reference point could be used to influence CDNI path selection using request attributes associated with ATIS C2 Extended Content Origin functionality such as the Scheduled Transmission Service. If the ATIS Service Provider requested the use of the ATIS C2 Scheduled Transmission Service via the ATIS C2 reference point, then the CDN selection path selected between the ATIS Content Provider and the ATIS Service Provider should select a CDN path capable of supporting

ATIS C2 Scheduled Transmission Service which is a component of ATIS C2 Extended Content Origin Functionality.

In this use case, the ATIS Service Provider CDN may takes a role in selecting the upstream CDN(s) on the path to the Content Service Provider's Origin Server. For example, the ATIS Service Provider CDN may require that the selected CDN path from the Content Service Provider to the ATIS Service Provider only include CDNs that support the ATIS C2 Extended Content Origin Functionality.

This use case matches the CDNI problem statement defined in [I-D.jenkins-cdni-problem-statement] but introduces a new requirement that the CDNI protocol suite must enable the discovery and selection of downstream CDNs based on additional attributes such as the ability of a CDN to support the ATIS C2 Extended Content Origin Functionality.

5. ATIS Interfaces to CDNI

Figure 5 illustrates how the ATIS architecture specificed in [ATIS-08000042] may interface with the CDNI protocols specified in [I-D.jenkins-cdni-problem-statement] within the content of the Internet Sourced Content and Off Net delivery use cases.

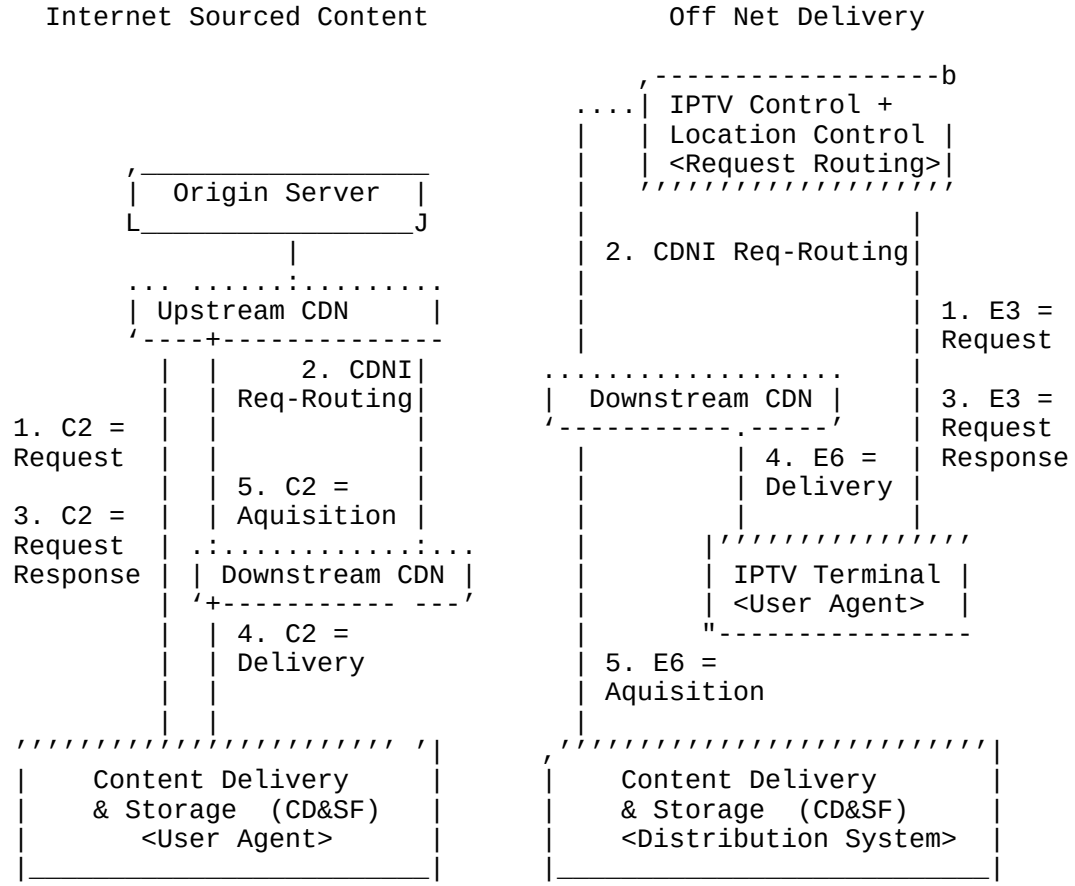

Figure 5: ATIS - CDNI Interfaces

The protocol sequences and interfaces used in the Off Net use case are illustrated on the right side of Figure 5 while the protocol sequences and interfaces used in the Internet Sourced Content use case are illustrated on the left side of Figure 5.

In the Off Net use case, the IPTV terminal acting as a CDNI User Agent inititates the CDNI Request Routing function by issuing an HTTP Get or RTSP Setup request via the E3 reference point (step 1). The ATIS Location Control Function uses relevant information from the E3 reference point (such as the IP address of the IPTV Terminal and the request URI) to issue a CDNI Request Routing request to the candidate Downstream CDNs (step 2). Once a Downstream CDN is selected, the IPTV Terminal is provided with a reference to the selected Downstream CDN via an E3 redirect response (step 3). Once the redirect response is received, the IPTV Terminal re-issues the HTTP Get or RTSP Setup

request via reference point E6 which in this case acts as a CDNI Delivery interface (step 4). On a cache miss, the Downstream CDN forwards the HTTP Get or RTSP Setup request to the Content Delivery and Storage Function which responds with the requested content (step 5).

In the Internet Sourced Content use case, the Content Delivery and Storage Function acting as a CDNI User Agent initiates the CDNI Request Routing function by issuing an HTTP Get via the C2 reference point (step 1). The Upstream CDN uses the relevant information from the C2 reference point (such as the IP address of the requesting CD&SF and the request URI) to issue a CDNI Request Routing request to the candidate Downstream CDNs (step 2). Once a Downstream CDN is selected, the CD&SF is provided with a reference to the selected Downstream CDN via an C2 redirect response (step 3). Once the redirect response is received, the CD&SF re-issues the HTTP Get request via reference point C2 which in this case acts as a CDNI Delivery interface (step 4). On a cache miss, the Downstream CDN forwards the HTTP Get request to the Upstream CDN which responds with the requested content (step 5).

6. IANA Considerations

This document makes no request of IANA.

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

7. Security Considerations

To be added.

8. Acknowledgements

We thank Francois Le Faucheur for his guidance and contribution for developing this document.

9. References

9.1. Normative References

[I-D.bertrand-cdni-use-cases]
Bertrand, G., Stephan, E., Watson, G., Burbridge, T., and P. Eardley, "Use Cases for Content Distribution Network

Interconnection", draft-bertrand-cdni-use-cases-01 (work in progress), January 2011.

[I-D.jenkins-cdni-problem-statement]
Niven-Jenkins, B., Faucheur, F., and N. Bitar, "Content Distribution Network Interconnection (CDNI) Problem Statement", draft-jenkins-cdni-problem-statement-02 (work in progress), March 2011.

[I-D.lefaucheur-cdni-requirements]
Faucheur, F., Viveganandhan, M., Watson, G., and Y. Lee, "Content Distribution Network Interconnection (CDNI) Requirements", draft-lefaucheur-cdni-requirements-01 (work in progress), March 2011.

[I-D.watson-cdni-use-cases]
Watson, G., "CDN Interconnect Use Cases", draft-watson-cdni-use-cases-00 (work in progress), January 2011.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[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.

9.2. Informative References

[I-D.gilletti-cdnp-aaa-reqs]
"CDI AAA Requirements, draft-gilletti-cdnp-aaa-reqs-01.txt", June 2001.

[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.

[ATIS-0800042]
ATIS IPTV Content on Demand Service
December 8, 2010

Authors' Addresses

Bruce Thompson
Cisco Systems
520 W Tasman Dr
San Jose, CA 95134
USA

Phone: +1 408 527 0446
Email: brucet@cisco.com

Akira Kobayashi
NEC
Tokyo
Japan

Phone: +81-44-455-8362
Email: a-kobayasi@ce.jp.nec.com