Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: July 21, 2012

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

January 18, 2012

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

Abstract

Content Delivery Networks (CDNs) are commonly used for improving 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 focuses on use cases that correspond to identified industry needs and that are expected to be realized once a CDN Interconnection (CDNI) solution is available. This document can be used to provide guidance to the CDNI WG about the interconnection arrangements to be supported and to validate the requirements of the various CDNI interfaces.

Table of Contents

1. Introduction

Content Delivery Networks (CDNs) are commonly used for improving 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 focuses on use cases that correspond to identified industry needs and that are expected to be realized once a CDNI solution is available. This document can be used to provide guidance to the CDNI WG about the interconnection arrangements to be supported and to validate the requirements of the various CDNI interfaces.

This document identifies the 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)

Then, the document highlights the need for interoperability to exchange and enforce content delivery policies (Section 5).

1.1. Terminology

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

We extend this terminology with the following terms.

Access CDN:

A CDN that is directly connected to the End User's access. An Access CDN may have specific information about the End User and the network, for instance, End User's profile and access capabilities.

Delivering CDN:

The CDN that delivers the requested piece of content to the End User. In particular, the Delivering CDN can be an Access CDN.

1.2. Abbreviations

- CDN: Content Delivery Network also known as Content Distribution Network

- CSP: Content Service Provider
- dCDN: downstream CDN
- DNS: Domain Name System
- DRM: Digital Rights Management
- EU: End User
- ISP: Internet Service Provider
- NSP: Network Service Provider
- QoE: Quality of Experience
- QoS: Quality of Service
- uCDN: upstream CDN
- URL: Uniform Resource Locator
- WiFi: Wireless Fidelity

1.3. Rationale 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 (decreased round-trip-time between the user and the delivery server) and better robustness,
- reduce the network operator's costs; for instance, lower delivery cost (reduced bandwidth usage) for cacheable content,
- reduce the Content Service Provider's (CSP) costs, such as datacenter capacity, space, and electricity consumption, as popular content is delivered through the CDN rather than through the CSP's servers.

Indeed, many Network Service Providers (NSPs) 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.

An example is depicted in Figure 1. Two CDN Providers establish a CDN Interconnection. 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 requests content from CSP-1, CDN-A considers that delivery by CDN-B is appropriate, for instance, because CDN-B is an Access CDN and the user is directly attached to it. CDN-A has delegated the handling of requests for CSP-1's content through the CDN Interconnection agreement, thus, the content is actually delivered from CDN-B.

The End User benefits from this arrangement through a better Quality of Experience (QoE), because the content is delivered from a nearby Surrogate. CDN Provider 'A' benefits because it does not need to deploy such an extensive CDN, whilst CDN Provider 'B' may receive 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 CSP-1 had also gone to the trouble of making a business agreement with CDN Provider 'B'.

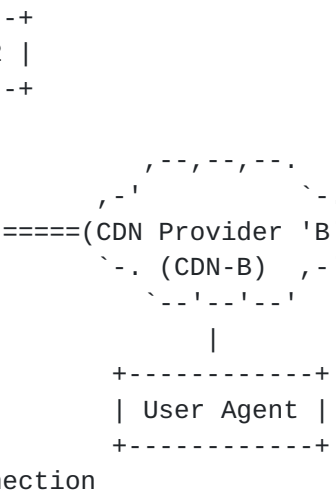

Figure 1

To extend the example, another Content Service Provider, CSP-2, 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-2 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.

1.4. The Need for CDN Interconnection Standards

The problem statement draft [I-D.ietf-cdni-problem-statement] describes extensively the CDNI problem space and explains why CDNI standards are required.

Existing CDN interfaces are proprietary and have often been designed for intra-CDN/intra-domain operations. Consequently, an external CDN typically cannot use these interfaces, especially if the two CDNs to be interconnected rely on different implementations. Nevertheless, [I-D.bertrand-cdni-experiments] shows that some level of CDN Interconnection can be achieved experimentally without standardized interfaces between the CDNs. However, 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 IETF CDNI WG's solution is, therefore, to overcome such shortcomings; a full list of requirements is being developed in [I-D.ietf-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 to its CSPs:

- without compromising the quality of delivery,
- without incurring additional transit and other network costs that would result from serving content 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 these 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 France and various countries in 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. Inter-Affiliates 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. For example, the CDN Provider might want to expose a single set of interfaces to the CSPs.

2.3. ISP Handling of Third-Party Content

Consider an ISP carrying to its subscribers a lot of content that comes from a third party CSP and that is injected into the access network by an Authoritative CDN Provider. There are mutual benefits to the Access CDN, the Authoritative CDN, and the CSP that would make a case for establishing a CDNI agreement. For example:

- Allow the CSP to offer improved QoE and QoE services to subscribers, for example, QoS and reduced round trip time.
- Allow the Authoritative CDN to reduce hardware capacity and footprint, by using the ISP caching and delivery capacity.
- Allow the ISP to reduce traffic load on some segments of the network by caching inside of the ISP network.
- Allow the ISP to influence and/or control the traffic ingestion points.
- Allow the ISP to derive some incremental revenue for transport of the traffic and to monetize QoE services.

2.4. Nomadic Users

In this scenario, a CSP wishes to allow End Users who move between CDNs to continue to access their content. The motivation of this case is to allow nomadic End Users to maintain access to content with a consistent QoE, across a range of devices and/or geographic

regions.

This use case covers situations like:

- End Users moving between different CDN Providers, which may reside within the same geographic region or different geographic regions,

- End Users switching between different devices or delivery technologies, as discussed in Section 4.

The term "Nomadic" does not necessarily relate to geographic roaming.

Consider the following example, illustrated in Figure 2: End User A has subscription to a broadband service from NSP A, her "home NSP". NSP A hosts CDN-A. Ordinarily, when End User A accesses content via NSP A (her "home NSP") the content is delivered from CDN-A, which in this example is within NSP A's network.

However, while End User A is not connected to NSP A's network, for example, because it is connected to a WiFi provider or mobile network, End User A can also access the same content. In this case, End User A may benefit from accessing the same content but delivered by an alternate CDN (CDN-B), in this case, hosted in the network of the WiFi or mobile provider, rather than from CDN-A in NSP A's network.

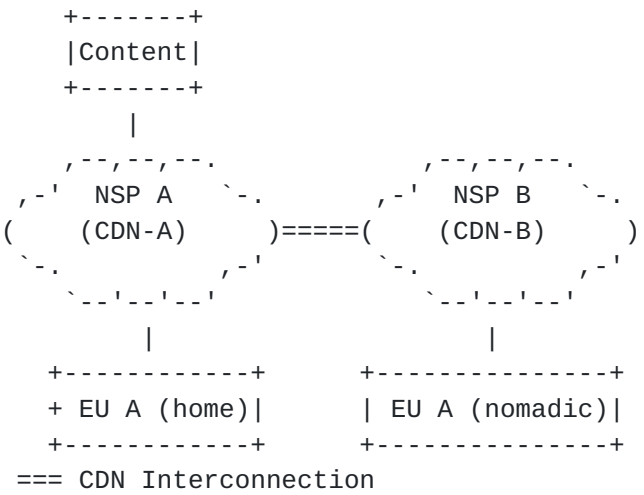

Figure 2

The alternate CDN (CDN-B) is allowed to distribute the content of CSP A to End User A; however, no other End Users in the region of CDN B are allowed to retrieve the content unless they too have such an

agreement for nomadic access to content.

Depending on CSP's content delivery policies (see Section 5.1), a user moving to a different geographic region may be subject to geo-blocking content delivery restrictions. In this case, he/she may not be allowed to access some pieces of content.

3. Offload Use Cases

3.1. Overload Handling and Dimensioning

A CDN is likely to be dimensioned to support an expected maximum traffic load. However, unexpected spikes in content popularity (flash crowd) 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 respective 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 competition. 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.

In this use case, the Delivering CDN on which requests are offloaded should be able to handle the offloaded requests. Therefore, the uCDN might require information on the dCDNs to be aware of the amount of traffic it can offload to every dCDN.

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 has at least two options: (1) depending on traffic management policies, forward some requests to the CSP's origin servers, and (2) redirect some requests toward

another CDN, which must be able to serve the redirected requests. The second option is a use case for CDNI.

3.2.2. Content Acquisition Resiliency

Source content acquisition may be handled in one of two ways:

- CSP origin, where a CDN acquires content directly from the CSP's origin server, or
- CDN origin, where a downstream CDN acquires content from a Surrogate within an upstream CDN.

The ability to support content acquisition resiliency, is an important use case for interconnected CDNs. When the content acquisition source fails, the CDN might switch to another content acquisition source. Similarly, when several content acquisition sources are available, a CDN might balance the load between these multiple sources.

Though other server and/or DNS load balancing techniques may be employed in the network, interconnected CDNs may have a better understanding of origin server availability and be better equipped to both distribute load between origin servers and attempt content acquisition from alternate origin servers when acquisition failures occur. When normal content acquisition fails, a CDN may need to try other origin server options, e.g.:

- an upstream CDN may acquire content from an alternate CSP origin server,
- a downstream CDN may acquire content from an alternate Surrogate within an upstream CDN, as directed by the upstream CDN's request routing interface,
- a downstream CDN may acquire content from an alternate upstream CDN, or
- a downstream CDN may acquire content directly from the CSP's origin server.

Though content acquisition protocols are beyond the scope of CDNI, the selection of content acquisition sources should be considered.

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 may wish to extend the supported range of devices and User Agents or the supported range of delivery technologies. In this case, a CDN Provider may interconnect with a CDN that offers services:

- that the CDN Provider is not willing to provide or,
- that its own CDN is not able to support

The following examples illustrate this use case:

1. CDN-A cannot support a specific delivery protocol. For instance, CDN-A may interconnect with CDN-B to serve a proportion of its traffic that requires HTTPS. CDN-A may use CDN-B's footprint (which may overlap with its own) to deliver HTTPS without needing to deploy its own infrastructure. This case could also be true of other formats, delivery protocols (RTMP, RTSP, etc.) and features (specific forms of authorization such as tokens, per session encryption, etc.).

2. CDN-A has footprint covering traditional fixed line broadband and wants to extend coverage to mobile devices. In this case, CDN-A may contract and interconnect with CDN-B who has both:

    * physical footprint inside the mobile network,
    * the ability to deliver content over a protocol that is required by specific mobile devices.

These cases can apply to many CDN features that a given CDN Provider may not be able to support or not be willing to invest in, and thus, that the CDN Provider would delegate to another CDN.

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. In addition, a CDN Provider may have a multi-vendor strategy for its CDN deployment. Finally, 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 requirements. As a result, the CDN Provider may establish a CDN Interconnection agreement with another CDN Provider that can provide the expected QoE to the End User, e.g., via an Access CDN able to deliver content from Surrogates located closer to the End User and with the required service level.

5. Enforcement of Content Delivery Policy

CSPs commonly require the ability to place delivery restriction on sets of content, which are provided by existing CDNs. The ability to support such delivery restrictions across interconnected CDNs is desirable, but depends on the capabilities of the involved CDNs. Thus, it is important to be able to detect and define when these features cannot be enforced.

5.1. Content Delivery Restrictions

The content distribution policies that a CSP attaches to a piece of content depend on many criteria. For instance, distribution policies for audiovisual content often combine:

- temporal constraints (e.g., available for 24 hours, available 28 days after DVD release, etc.),
- resolution-based constraints (e.g., high definition vs. standard definition), and
- geolocation-based constraints (e.g., per country).

CSPs may require from their CDN Providers that they translate some of the above requirements into content delivery policies for their CDNs. For instance, CDNs might implement "geo-blocking" rules specifying:

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

Similarly, an uCDN might implement some temporal constraints on content availability. For example, it could restrict access to pre-positioned content prior to the opening of the availability window or

disable the delivery of content from the dCDNs (e.g., through purging) after the availability window has closed.

5.2. Secure Access

Many protocols exist for delivering content to End Users. CSPs may often wish to dictate a specific protocol or set of protocols which are acceptable for delivery of their content, especially in the case where content protection or user authentication is required (e.g., must use HTTPS). CSPs may also wish to perform per-request authentication/authorization decision and then have the CDNs enforce that decision (e.g., must validate URL signing, etc.).

An uCDN needs to be able to exclude dCDNs which lack support for the secure access features requested by the CSP.

5.3. Branding

Preserving the branding of the CSP throughout delivery is often important to the CSP. CSPs may desire to offer content services under their own name, even when the associated CDN service involves other CDN Providers. For instance, a CSP may desire to ensure that content is delivered with URIs appearing to the endusers under the CSP's own domain name, even when the content delivery involves separate CDN Providers. The CSP may wish to forbid the delivery of its content by specific dCDNs that lack support for such branding preservation features.

Similar restrictions may exist when the uCDN wants to offer CDN services under its own branding even if dCDNs are involved. Conversely, a CDN Provider might not want the brand of a CDN Exchange to be visible, even if the CDN Exchange is involved in the content delivery call flow.

6. Acknowledgments

The authors would like to thank Kent Leung, Francois Le Faucheur, Ben Niven-Jenkins, and Scott Wainner for lively discussions, as well as for their reviews and comments on the mailing list.

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

7. IANA Considerations

This memo includes no request to IANA.

8. Security Considerations

CDN Interconnection, as described in this document, has a wide variety of security issues that should be considered.

In addition to the security considerations within the uCDN and dCDN, four contexts involve security and trust issues:

a. The relationship between the CSP and the uCDN: the main contract arrangement for distribution, which authorizes the uCDN to acquire content on CSP's origin servers and to deliver it, potentially with content delivery restrictions.

b. The relationship between the uCDN and dCDN: the transitive trust relationship that extends the contract defined in (a) above and that authorizes the dCDN to acquire content on uCDN's or CSP's origin servers and to deliver it, potentially with content delivery restrictions.

c. The relationship between the End User and dCDN: the recognition of right to download predicated on (b) above.

d. The relationship between the End User and the CSP: the contract that authorizes the End User to access to the content.

CDNI should enable the four parties (CSP, uCDN, dCDN, End User) to negotiate a security method or a method for confirming authorization along the chain of trust (CSP -> uCDN -> dCDN -> End User).

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 situations 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. CDNI interfaces and protocols should attempt to minimize overhead for dCDNs.

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

9. References

9.1. Normative References

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

9.2. Informative References

[I-D.bertrand-cdni-experiments]
Bertrand, G., Faucheur, F., and L. Peterson, "Content Distribution Network Interconnection (CDNI) Experiments", draft-bertrand-cdni-experiments-01 (work in progress), August 2011.

[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-02 (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.

[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 (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

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

Email: emile.stephan@orange.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