Use Cases for Content Delivery Network Interconnection

Abstract

   Content Delivery Networks (CDNs) are commonly used for improving the
   End User experience of a content delivery service while keeping cost
   at a reasonable level. This document focuses on use cases that
   correspond to identified industry needs and that are expected to be
   realized once open interfaces and protocols supporting the
   interconnection of CDNs are specified and implemented. This document
   can be used to motivate the definition of the requirements to be
   supported by CDN Interconnection (CDNI) interfaces. It obsoletes RFC
   3570.

Copyright Notice

Table of Contents

1. Introduction

   Content Delivery Networks (CDNs) are commonly used for improving the
   End User experience of a content delivery service while keeping cost
   at a reasonable level. This document focuses on use cases that
   correspond to identified industry needs and that are expected to be
   realized once open interfaces and protocols supporting the
   interconnection of CDNs are specified and implemented. The document
   can be used to motivate the definition of the requirements (as
   documented in [CDNI-REQ]) to be supported by the set of CDN
   Interconnection (CDNI) interfaces defined in [RFC6707].

   [RFC3570] describes slightly different terminologies and models for
   "Content Internetworking (CDI)". This document obsoletes RFC 3570 to
   avoid confusion.

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

   o CDN Footprint Extension Use Cases (Section 2)

   o CDN Offload Use Cases (Section 3)

   o CDN Capability Use Cases (Section 4)

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

1.1. Terminology

   In this document, the first letter of each CDNI-specific term is
   capitalized. We adopt the terminology described in [RFC6707].

   We extend this terminology with the following term:

   Access CDN:

   A CDN that includes Surrogates in the same administrative network as
   the End User. Such a CDN can use accurate information on the End
   User's network context to provide additional Content Delivery
   Services to Content Service Providers.

1.2. Abbreviations

   o CDN: Content Delivery Network, also known as Content Distribution
      Network

   o CSP: Content Service Provider

o dCDN: downstream CDN

o DNS: Domain Name System

o EU: End User

o ISP: Internet Service Provider

o NSP: Network Service Provider

o QoE: Quality of Experience

o QoS: Quality of Service

o uCDN: upstream CDN

o URL: Uniform Resource Locator

o WiFi: Wireless local area network (WLAN) based on IEEE 802.11

1.3. Rationale for CDN Interconnection

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

   o improve the experience for the End User; for instance delivery has
      lower latency (decreased round-trip-time and higher throughput
      between the user and the delivery server) and better robustness
      (ability to use multiple delivery servers),

   o reduce the network operator's costs; for instance, lower delivery
      cost (reduced bandwidth usage) for cacheable content,

   o reduce the Content Service Provider's (CSP) internal
      infrastructure costs, such as data center capacity, space, and
      electricity consumption, as popular content is delivered
      externally through the CDN rather than through the CSP's own
      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
   stand-alone 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, where 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.
Independently, CDN Provider 'A' and CDN Provider 'B' agree to
interconnect their CDNs.

When a given User Agent requests content from CSP-1, CDN-A may
consider that delivery by CDN-B is appropriate, for instance, because
CDN-B is an Access CDN and the user is directly attached to it.
Through the CDN Interconnection arrangements put in place between
CDN-A and CDN-B (as a result of the CDN Interconnection agreement
established between CDN Provider 'A' and CDN Provider 'B'), CDN-A can
redirect the request to CDN-B and the content is actually delivered
to the User Agent by CDN-B.

The End User benefits from this arrangement through a better Quality
of Experience (QoE, see [RFC6390]), because the content is delivered
from a nearby Surrogate (e.g., lower latency, bottlenecks avoided).
CDN Provider 'A' benefits because it does not need to deploy such an
extensive CDN, while CDN Provider 'B' may receive some compensation
for the delivery. CSP-1 benefits because it only needs to make one
business agreement and one technical arrangement 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 and technical
arrangement with CDN Provider 'B'.

```
   +-------+ +-------+
   | CSP-1 | | CSP-2 |
   +-------+ +-------+
       |         |
     ,--,--,--./          ,--,--,--.
   ,-'          `-.      ,-'          `-.
  (CDN Provider 'A')=====(CDN Provider 'B')
   `-.  (CDN-A) ,-'       `-. (CDN-B)  ,-'
      `--'--'--'              `--'--'--'
                                  |
                            +----------+
                            | End User |
                            +----------+
   === CDN Interconnection
```

                              Figure 1
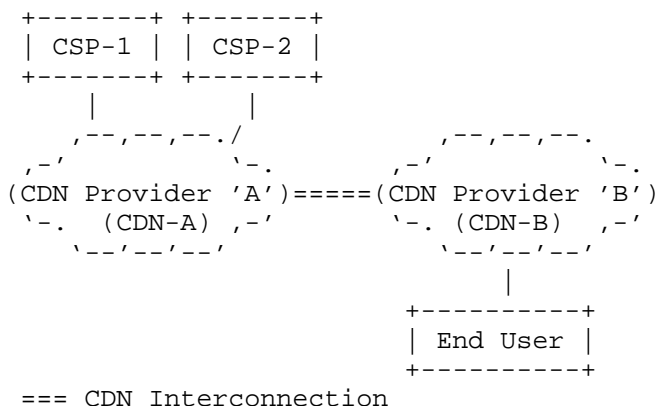
To extend the example, another Content Service Provider, CSP-2, may
also reach an agreement with CDN Provider 'A'. However, CSP-2 may
not want its content to be distributed by CDN Provider B; for
example, CSP-2 may not want to distribute its content in the area
where CDN Provider 'B' operates. This example illustrates that
policy considerations are an important part of CDNI.

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:

   o without compromising the quality of delivery.

   o without incurring additional transit and other network costs that
      would result from serving content from geographically or
      topologically remote Surrogates.

   o without incurring the cost of deploying and operating Surrogates
      and the associated CDN infrastructure that may not be justified in
      the corresponding geographic region (e.g., because of relatively
      low delivery volume, or conversely because of the high investments
      that would be needed to satisfy the high volume).

   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

   The previous section describes the case of geographic extension
   between CDNs operated by different entities. A large CDN Provider
   may have several subsidiaries that each operate their own CDN (which
   may rely on different CDN technologies, see Section 4.2). In certain

   circumstances, the CDN Provider needs to make these CDNs interoperate
   to provide consistent service to its customers on the whole
   collective footprint.

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 ISP's
   network by an Authoritative CDN Provider. There are mutual benefits
   to the ISP (acting as an Access CDN), the Authoritative CDN, and the
   CSP that would make a case for establishing a CDNI agreement. For
   example:

   o allowing the CSP to offer improved QoE and QoE services to
      subscribers, for example, reduced content startup time or
      increased video quality and resolution of adaptive streaming
      content.

   o allowing the Authoritative CDN to reduce hardware capacity and
      footprint, by using the ISP caching and delivery capacity.

   o allowing the ISP to reduce traffic load on some segments of the
      network by caching inside of the ISP network.

   o allowing the ISP to influence and/or control the traffic ingress
      points.

   o allowing 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
   access networks 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:

   o End Users moving between different access networks, which may be
      located within the same geographic region or different geographic
      regions.

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

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

However, while End User A is not connected to ISP 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 delivered by
an alternate CDN (CDN-B), in this case, hosted in the network of the
WiFi or mobile provider (ISP B), rather than from CDN-A in ISP A's
network.

```
      +-------+
      |Content|
      +-------+
          |
       ,--,--,--.                   ,--,--,--.
     ,-'  ISP A   `-.           ,-'  ISP B    `-.
    (     (CDN-A)      )=====(      (CDN-B)     )
     `-.          ,-'          `-.          ,-'
       `--'--'--'                `--'--'--'
          |                          |
    +------------+          +---------------+
    + EU A (home)|          | EU A (nomadic)|
    +------------+          +---------------+
    === CDN Interconnection
```
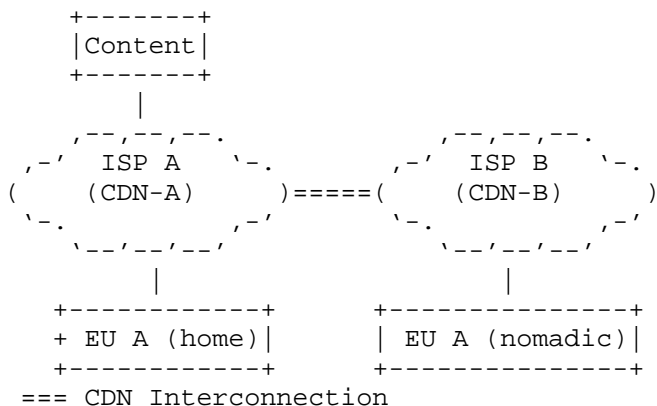
                             Figure 2

Though the content of CSP A is not accessible by typical End Users of
CDN-B, End User A is able to gain access to her "home" content (i.e.,
the content of CSP A) through the alternate CDN (CDN-B).

Depending on the CSP's content delivery policies (see Appendix A.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 each other's resources during their respective
peaks of activity.

Offload also applies to planned situations in which a CDN Provider
needs CDN capacity in a particular region during a short period of
time. For example, a CDN can offload traffic to another CDN for the
duration of a specific maintenance operation or for the distribution
of a special event, as in the scenario depicted in Figure 3. For
instance, consider a TV channel that is the distributor for a major
event, such as a celebrity's wedding or a major sport competition,
and this TV channel has contracted particular CDNs for the delivery.
The CDNs (CDN-A and CDN-B) that the TV channel uses for delivering
the content related to this event are likely to experience a flash
crowd during the event and will need to offload traffic, while other
CDNs (CDN-C) will support a more typical 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 each dCDN.

```
   +------------+
   | TV Channel |
   +------------+
      |         \
    ,-,--,-.     \ ,-,--,-.           ,-,--,-.
  ,'        `.   ,'        `.       ,' CDN-C  `.
 (   CDN-A    ) (   CDN-B    )==(   offload   )
  `.        ,'   `.        ,'     `.        ,'
    `-'--'-'        `-'--'-'         `-'--'-'


   === CDN Interconnection
```
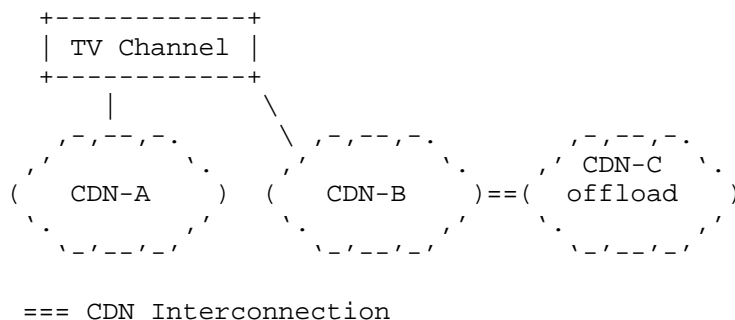
                            Figure 3

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 three options:

   1. if possible, use internal mechanisms to redirect traffic onto
       surviving equipment,

   2. depending on traffic management policies, forward some requests
       to the CSP's origin servers, and/or

   3. redirect some requests toward another CDN, which must be able to
       serve the redirected requests.

   The last option is a use case for CDNI.

3.2.2. Content Acquisition Resiliency

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

   o CSP origin, where a CDN acquires content directly from the CSP's
      origin server, or

   o 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 content sources when acquisition failures
   occur. When normal content acquisition fails, a CDN may need to try
   other content source options, for example:

   o an upstream CDN may acquire content from an alternate CSP origin
      server,

   o a downstream CDN may acquire content from an alternate Surrogate
      within an upstream CDN,

   o a downstream CDN may acquire content from an alternate upstream
      CDN, or

   o 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 and facilitated.

4. 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:

o the CDN Provider is not willing to provide, or

o 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 [RFC2818]. 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 (e.g., the Real Time Messaging Protocol (RTMP), the Real Time Streaming Protocol (RTSP), etc.), and features (specific forms of authorization such as tokens, per session encryption, etc.).

2. CDN-A has a 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:

    * a physical footprint inside the mobile network,

    * the ability to deliver content over a protocol that is required by specific mobile devices.

3. CDN-A only supports IPv4 within its infrastructure but wants to deliver content over IPv6. CDN-B supports both IPv4 and IPv6 within its infrastructure. CDN-A interconnects with CDN-B to serve out its content over native IPv6 connections.

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 would not 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 that is able to deliver content from Surrogates located
   closer to the End User and with the required service level.

5. Enforcement of Content Delivery Policy

   An important aspect common to all the above use cases is that CSPs
   typically want to enforce content delivery policies. A CSP may want
   to define content delivery policies that specify when, how, and/or to
   whom the CDN delivers content. These policies apply to all
   interconnected CDNs (uCDNs and dCDNs) in the same or similar way that
   a CSP can define content delivery policies for content delivered by a
   single, non-interconnected CDN. Appendix A provides examples of CSP-
   defined policies.

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.

   Finally, the authors acknowledge the work of the former CDI working
   group. This document obsoletes [RFC3570] to avoid confusion.

7. Security Considerations

   This document focuses on the motivational use cases for CDN
   Interconnection and does not analyze the associated threats. Those
   threats are discussed in [RFC6707]. Appendix A.2 of this document
   provides example security policies that CSPs might impose on CDNs to
   mitigate the threats.

8. References

8.1. Normative References

   [RFC6707]   Niven-Jenkins, B., Le Faucheur, F., and N. Bitar,
               "Content Distribution Network Interconnection (CDNI)
               Problem Statement", RFC 6707, September 2012.

8.2. Informative References

   [CDNI-REQ] Leung, K. and Y. Lee, "Content Distribution Network
               Interconnection (CDNI) Requirements", Work in Progress,
               June 2012.

   [RFC2818]   Rescorla, E., "HTTP Over TLS", RFC 2818, May 2000.

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

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

Appendix A. Content Service Providers' Delivery Policies

   CSPs commonly apply different delivery policies to given sets of
   content assets delivered through CDNs. Interconnected CDNs need to
   support these policies. This appendix presents examples of CSPs'
   delivery policies and their consequences on CDNI operations.

A.1. Content Delivery Policy Enforcement

   The content distribution policies that a CSP attaches to a content
   asset may depend on many criteria. For instance, distribution
   policies for audiovisual content often combine constraints of varying
   levels of complexity and sophistication, for example:

   o temporal constraints (e.g., available for 24 hours, available 28
      days after DVD release, etc.),

   o user agent platform constraints (e.g., mobile device platforms,
      desktop computer platforms, set-top-box platforms, etc.),

   o resolution-based constraints (e.g., high definition vs. standard
      definition encodings),

   o user agent identification or authorization,

   o access network constraints (e.g., per NSP), and

   o IP geo-blocking constraints (e.g., for a given coverage area).

   CSPs may use sophisticated policies in accordance with their business
   model. However, the enforcement of those policies does not
   necessarily require that the delivery network understand the policy
   rationales or how policies apply to specific content assets. Content
   delivery policies may be distilled into simple rules that can be
   commonly enforced across all dCDNs. These rules may influence dCDN
   delegation and Surrogate selection decisions, for instance, to ensure
   that the specific rules (e.g., time-window, geo-blocking, pre-
   authorization validation) can indeed be enforced by the Delivering
   CDN. In turn, this can guarantee to the CSP that content delivery
   policies are properly applied.

```
   +-----+
   | CSP |  Policies driven by business (e.g., available only
   +-----+  in the UK and only from July 1st to September 1st)
      \
       \ Translate policies into
        \simple rules (e.g., provide an authorization token)
         \
          V
        +-----+
        | CDN |  Apply simple rules (e.g., check an
        +-----+  authorization token and enforce geo-blocking)
           \
            \ Distribute simple rules
             V
          +-----+
          | CDN |  Apply simple rules
          +-----+
```
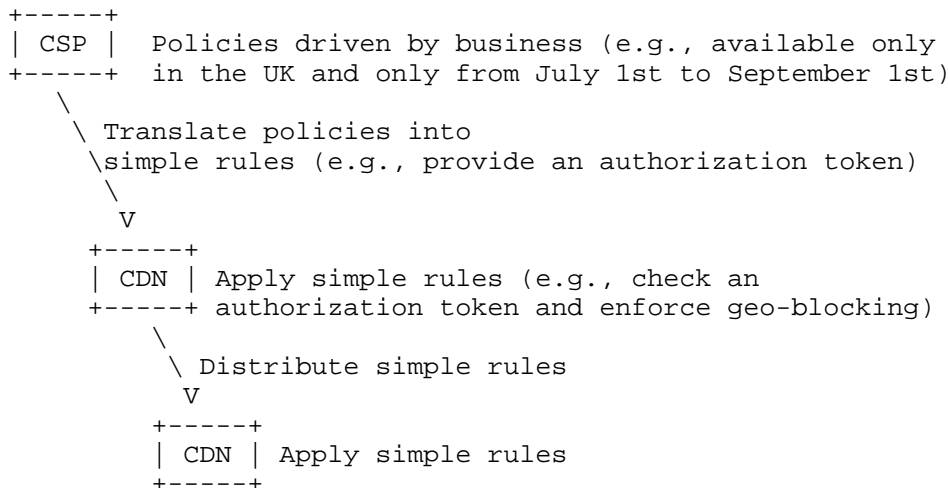
                              Figure 4

A.2. Secure Access

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

A.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 End Users under the
   CSP's own domain name, even when the content delivery involves
   separate CDN Providers. The CSP may wish to prevent the delivery of
   its content by specific dCDNs that lack support for such branding
   preservation features.

   Analogous cases exist when the uCDN wants to offer CDN services under
   its own branding even if dCDNs are involved, and so it restricts the
   delivery delegation to a chain that preserves its brand visibility.

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
   FR
   EMail: emile.stephan@orange.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 J. Ma
   Azuki Systems, Inc.
   43 Nagog Park
   Acton, MA 01720
   USA
   Phone: +1 978-844-5100
   EMail: kevin.ma@azukisystems.com

   Grant Watson
   Alcatel-Lucent (Velocix)
   3 Ely Road
   Milton, Cambridge CB24 6AA
   UK
   EMail: gwatson@velocix.com