Society of Cable Telecommunications Engineers

SCTE Cable-Tec Expo 2011

Motivation and Architecture for the Interconnection of Content Delivery Networks

Richard Woundy, SVP Software & Applications
Weidong Mao, Senior Fellow
Yiu L. Lee, Distinguished Engineer
Carl Williams, Consultant
Comcast Cable

Address: 1701 JFK Boulevard, Philadelphia, PA 19103
Email: richard_woundy@cable.comcast.com, weidong_mao@cable.comcast.com, carl_williams3@cable.comcast.com, yiu_lee@cable.comcast.com

Abstract

This paper describes the interconnection of CDNs which are deployed and operated by independent providers. Current practice is that CDNs are operated independently and autonomously, for example by cable operators to offer video on demand services to set-top boxes, or by Internet providers to enable over-the-top video downloads and streaming. This is in contrast to the proposed interconnection model, in which two independent CDN providers cooperate in the distribution of content assets from the origins to the end-users. While one CDN may ingest from the content origin, an end-user may download that same content from the edge node belonging to the other CDN.

CDN providers that are also broadband access providers may leverage interconnection in order to reduce backbone network traffic and costs, to increase CDN revenues, to extend their VOD services to other providers, and to enable higher quality content distribution to their broadband access customers. CDN providers may choose to offer services directly to content providers for their own network footprint, or to offload traffic from existing Internet CDN providers, or to organize a CDN federation with its own worldwide network footprint.

Basic request routing between CDN providers has been recently demonstrated by France Telecom/Orange, but additional architecture and specification work is needed. The CDN Interconnection (CDNI) group in the IETF is focused on agreement on specifications for the exchange of information for request routing, metadata information exchange, exchange of transaction logs and monitoring information, and the management of content (e.g. purging stale assets). Other standards bodies (ATIS and ETSI) are also working on CDN interconnection to support IPTV distribution and similar use cases.

Disclaimer: This technical paper explores various potential new technologies enabled for broadband service providers, but does <u>not</u> represent a discussion of Comcast business plans or priorities.

1. Introduction

Content Delivery Networks (CDNs) are commonly used for improving the footprint, network utilization, and the end-user experience of a content delivery service, at a reasonable cost. The way in which the CDNs do this is by routing an end user's request for content to the nearest super-pop or a location where the actual content is stored or cached. Following this, the requested content returns to the end user through a caching server physically located nearer to the end user. This has the effect of the first request having the standard lag of a long distance request, but subsequent requests being as fast as if the content were more regionally distributed. CDNs improve network performance by maximizing bandwidth, improving accessibility and maintaining correctness through content replication. In addition to the performance advantages, it also allows the content to be taken off line.

CDNs have become the standard delivery method for services such as on-demand and live delivery of video, music, games and social media to broadband and mobile users. Internet Video has evolved in different forms in the past 5 years. The simplest form is a low resolution small format image that is commonplace on almost all websites today. The next step up is video streaming or downloading sites offered by wide range of media outlets with content ranging from user generated content often encoded at low resolution – YouTube – to TV and movie clips, previews and trailers, often encoded at somewhat higher resolution. Cloud-based Over-The-Top (OTT) video streaming has become highly popular with significant increases in demand for delivery of feature-length movies, e.g., Netflix. The evolution doesn't stop with the characteristics of the video content, but with how the user obtains the content. With YouTube it is primary an ad-based service, while with Netflix it is subscription-based. iTunes is a fee based service per download.

In a recent study by Sandvine [7], entitled 2011Global Internet Phenomena, it was reported that Real-Time Entertainment is the largest contributor to data consumption on both fixed (43% of peak period traffic) and mobile access (41%) networks. Within this category the 2011 Sandvine report on Netflix traffic [8] found that more than 29% of downstream traffic in the US is from consumers downloading movies and video content from Netflix and was heaviest between 8pm-10pm. This success of Internet video, however, can be a mixed blessing for broadband service providers. Increased use of Internet video threatens to drive up Internet traffic and operation costs but does nothing for revenue. The dynamic changes of how video is seen, delivered and charged for by the content service providers is driving new business models based on a more open architecture for delivery of these services.

As one example, Netflix pays multiple CDN providers to do the streaming for them. These CDN providers work in proprietary fashion and independently of each other by reproducing the same functions across its delivery network. The footprint of any of these CDNs may not extend close enough to the end user's location to realize the cost benefits and superior user experience that a more distributed CDN would provide. In addition, CDN providers may wish to use multiple vendor solutions for implementation. It was found in the France Telecom experiments [3] that interconnection of CDN solutions exposed gaps and provided a basis for IETF standardization work for CDN interconnection. The IETF work is focused more on what is minimally required to interconnect cooperating CDNs than it is to explore the entire scope of CDNI. In addition ATIS and ETSI are both working on CDN Interconnection use cases and requirements documents. The goal of the ATIS, IETF and ETSI CDNI efforts is to facilitate the delivery of interoperable, secure, and managed services, and to enable seamless, global content delivery, especially of video, between independent CDNs.

2. Abbreviations and Acronyms

ATIS	Alliance for Telecommunications Industry Solutions
CDN	Content Delivery Network
CDNI	Content Delivery Network Interconnection
CSF	Cloud Services Forum
CSP	Content Service Provider
dCDN	Downstream CDN
ETSI	European Telecommunications Standards Institute
IETF	Internet Engineering Task Force
IIF	IPV Interoperability Forum
IPTV	Internet Protocol Television
IPV	Internet Protocol Video
ISP	Internet Service Provider
MCD	Media Content Distribution
OTT	Over The Top
uCDN	Upstream CDN
TISPAN	Telecommunications and Internet converged Services and Protocols for Advanced Networking

3. CDN Interconnect Business Case

There are two key business motivations for ISP involvement in content delivery. The first motivation is to enable efficient distribution of ISP-licensed content to customers; one example of such a system would be the Comcast Content Delivery Network, used for its video-on-demand platform [9]. The second motivation is to optimize the delivery of third party content over the ISP's network, such as the distribution of Netflix or YouTube content.

CDN Interconnection (CDNI) benefits ISP-licensed content distribution in several ways. First, an ISP with its own CDN may use CDNI to leverage an external CDN for delivery of content to subscribers that are beyond the footprint of the ISP. Second, a larger ISP and a smaller ISP may partner through CDN interconnection, such that the larger ISP bears the responsibilities of ingesting, transcoding, and packaging licensed content, and the smaller ISP focuses on caching licensed content to minimize network costs for both parties.

CDNI may be beneficial for the delivery of third party content over the Internet as well. By caching third party content in the ISP network, the ISP and the Content Service Provider (CSP) may reduce network traffic, minimizing costs for both parties. The ISP may generate additional revenue for this caching service, since this service reduces the costs of the upstream CDN provider. Since ISP content caches would likely be located closer to the content consumers, network latency may be reduced and network reliability may be improved, leading to an improved experience for end users.

4. Architecture

4.1. Traditional Content Delivery Network Overview

For traditional OTT content delivery, when the CSP receives a request from the user agent, it will choose one of its contracted CDN providers to deliver the content. The CDN provider selection process is often based on the information in the user's request and the user's location. The ISP providing Internet service to the user is unaware of the selection process. From the ISP perspective, its responsibility is to deliver the IP packets from the CDN provider to the user agent. Even when multiple user agents in the same ISP request the same content, the CDN provider will deliver the same content multiple times to the ISP. Figure 1 shows the high-level architecture.

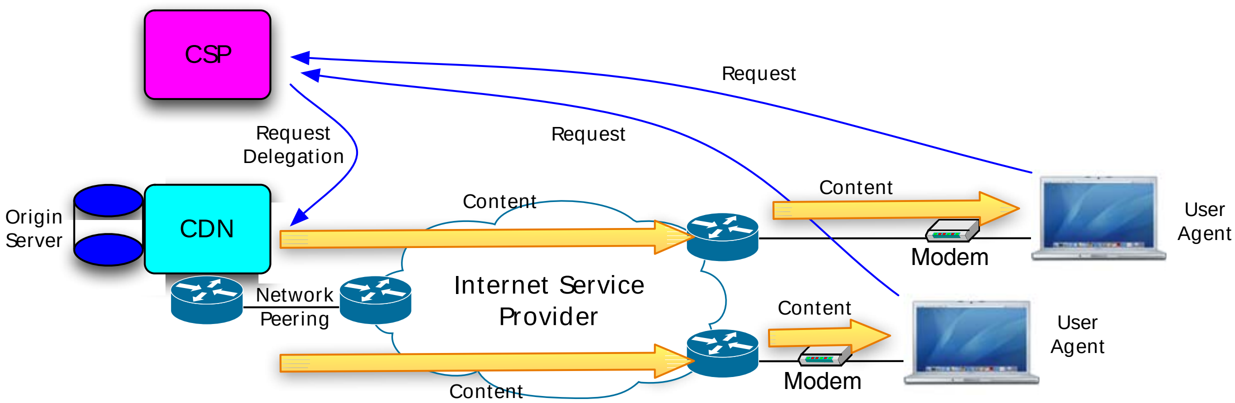

Figure 1 – Illustration of CDN Architecture

In Figure 1, the CDN provider receives two requests from the CSP to deliver the same content to two user agents. The CDN provider would deliver the same content twice to the ISP network. The ISP is unaware of the content delivery service, nor does it know anything about the content.

This delivery method has two shortcomings. First, both the CDN provider and the ISP would need to build an infrastructure that can sustain the peak-time traffic. Consider that most popular content will likely be requested multiple times during peak hours, thus the same piece of content will need to be delivered multiple times. This is inefficient and expensive. Second, the CDN provider does not have visibility to the ISP's network topology. It may choose a peering point that is not optimal to deliver the content to the user. This may degrade the service due to network latency and other networking factors.

4.2. Content Delivery Network Interconnection Overview

Alternatively, the CDN provider can peer with the ISP's CDN to deliver the content via CDN Interconnection. When a user agent requests content, the CDN provider will determine where the request comes from, then delegate the request to the user's ISP CDN. If the ISP CDN has not yet cached the content, it would obtain the content from the CDN provider and cache it in its own CDN. When the next user requests the same content, the CDN provider can redirect the request to the ISP CDN. The ISP CDN would deliver the content from its cache rather than asking the CDN provider to deliver it. This is the basic idea of CDN-Interconnect (CDNI) architecture. Figure 2 shows the high-level architecture.

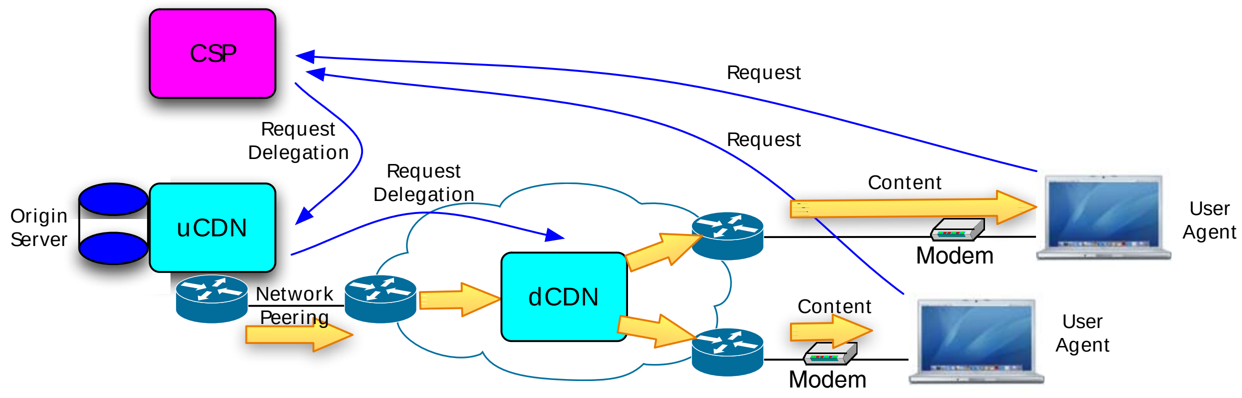

Figure 2 – Illustration of CDN-Interconnect (CDNI) Architecture

In Figure 2, the CDN provider (Upstream CDN or uCDN) delegates the two user's requests to the ISP provider (Downstream CDN or dCDN). If the ISP CDN does not have the content, it would ask the uCDN for the location of the content and obtain it. If the content has already been cached, the dCDN would directly deliver it to the user instead of the CDN provider delivering it.

CDNI would address the two shortcomings mentioned in Section 4.1. First, the uCDN only needs to stream the content once to the dCDN; this will help the uCDN to better utilize its network capacity. The ISP can deploy the dCDN at the edge of the network so that it can also save network capacity in its core network. Second, the ISP knows its network topology, so it can select the most effective route to deliver the content to its users.

CDNI is not limited to delegation from the CDN provider to the ISP CDN. In reverse, an ISP can delegate a request to a CDN provider. For example: an ISP subscriber roaming outside his ISP's service area may want to access some local content provided by his ISP. The ISP may delegate the request to a CDN provider which has presence in the subscriber's current location. Furthermore, a CDN provider can also peer with another CDN provider. For example: CDN-A in the USA may want to delegate the request to CDN-B in the UK, when a user in the UK requests content. All these peering scenarios are possible in the CDNI architecture, enabling a wide variety of business arrangements.

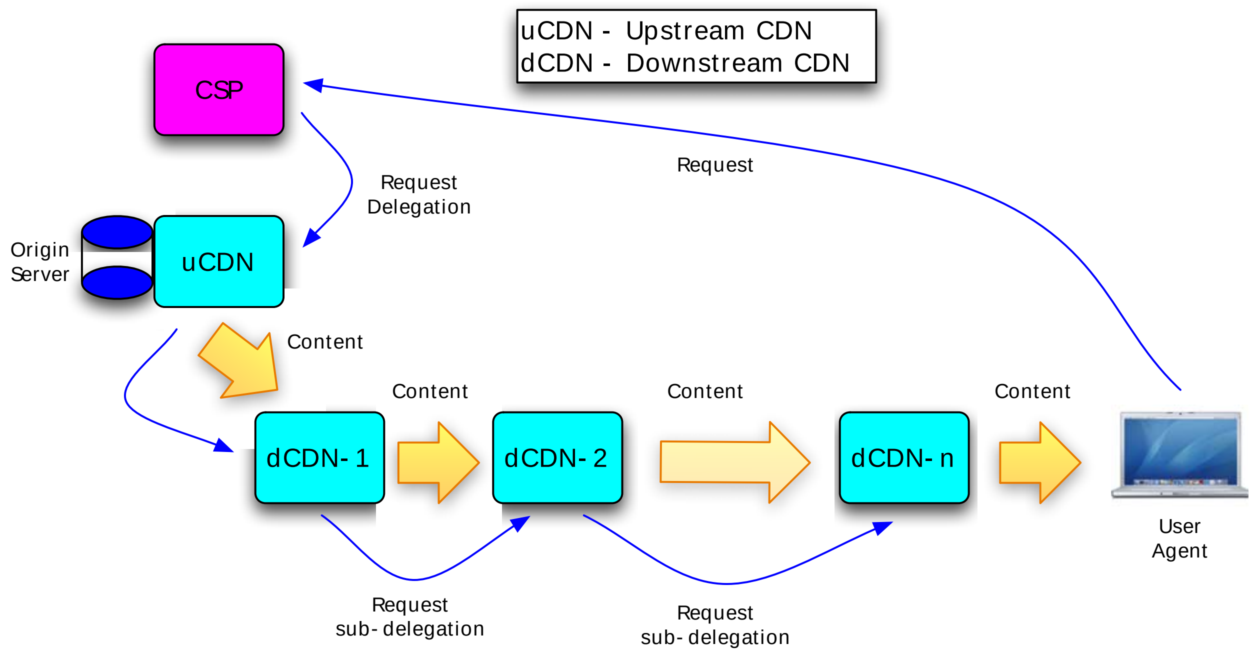

Figure 3 – Illustration of Multi-level Delegation with CDN Interconnection

In Figure 3, uCDN delegates the request to dCDN-1. Then, dCDN-1 delegates the request to dCDN-2. This delegation can continue through an arbitrary number of levels (i.e. *n-level*). In reality, this process can't go on forever because the user's request will eventually timeout. Figure 3 illustrates that the architecture allows multi-level delegation.

4.3. Content Delivery Network Internal Architecture

A typical CDN architecture includes a control plane and a set of edge caches[1].The control plane manages CDN resources, handles user requests, selects edge caches for service delivery, enables event logging, and distributes the metadata to the edge caches. The edge cache is the edge device which stores the cached content and delivers the content to users. Each content file normally is associated with a single metadata file. However, a metadata file can be associated with more than one content item. The metadata file contains policies such as serving region, availability windows, and expiration time. Figure 4 shows the architecture.

[1] We use Edge Cache in this document to represent the caching node serving the end users. In some documents (e.g. IETF), Surrogate is used to refer to the same function.

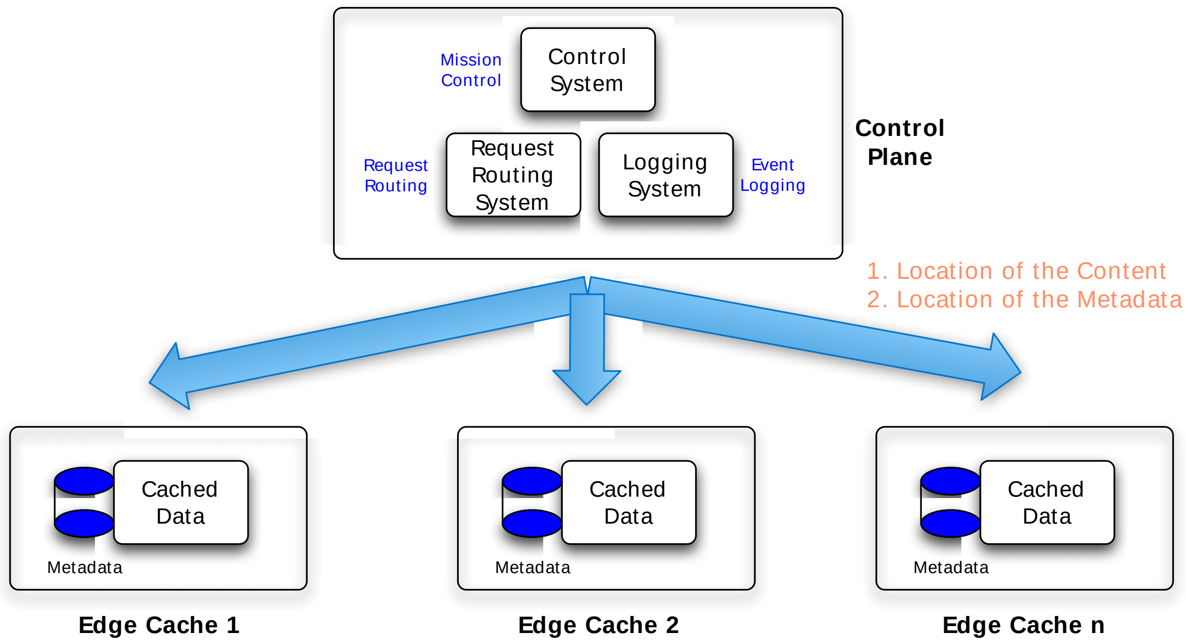

Figure 4 – Illustration of CDN Internal Architecture

1. The Control System is the "mission control" of the CDN. Its functions include (1) deciding whether to pre-position or to dynamically acquire content, (2) maintaining the location of the content and metadata, (3) monitoring the health of the edge cache and taking failed edge caches out of service, and (4) ensuring that the edge caches purge a piece of content before its expiration time.

2. The Request Routing System handles user requests. It contains a set of rules and policies to select the best edge cache for the user. Typically, the Request Routing system uses either the DNS or HTTP protocol to redirect a user to the best edge cache to get the content.

3. The Logging System collects events from the CDN system such as user requests, content delivery status, and other CDN related information.

4.4. Content Delivery Network Interconnect Architecture

CDNI focuses on the inter-domain CDN interaction. The uCDN and dCDN must interface to exchange enough information to complete the request. They also need to exchange logging information for billing settlement and trouble-shooting. In CDNI, the dCDN's edge cache will deliver the content, so the dCDN's edge cache must have the associated metadata of the content to govern the delivery. Figure 5 shows the necessary interfaces for CDNI.

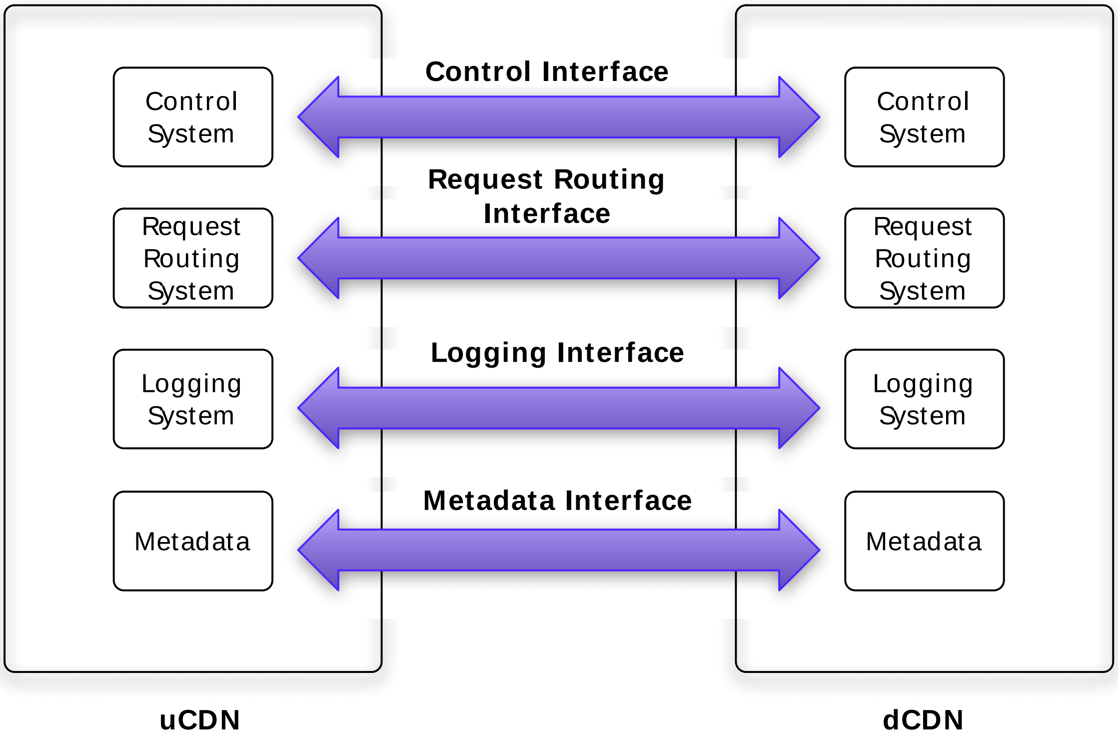

Figure 5 – Illustration of CDNI Interfaces

1. The Control System Interface is used to bootstrap the interconnection between two CDN systems. It enables two CDN systems to exchange information such as capacity, availability, system status and peering policies. It also allows the uCDN to purge any content in the dCDN at any given time.

2. The Request Routing Interface is used to support delegation of user requests from the uCDN to the dCDN. The dCDN can use this interface to report back to the uCDN the status of the request delegation. This interface also allows the uCDN to query the dCDN information before delegating the request.

3. The Logging Interface is used to collect events from dCDN to uCDN. This information is particularly important for billing, settlement, and troubleshooting. The uCDN uses the information collected via this interface to reconstruct the end-to-end service delivery of a request.

4. The Metadata Interface is used by uCDN to deliver the CDNI metadata of the content cached in the dCDN. The CDNI Metadata is specific to CDNI operation, and is distinct from other content metadata such as artist information and/or video encoding information. CDNI Metadata would include availability windows, geo-blocking information, distribution policy and content management information. This information is important for the uCDN to control the content distributed by the dCDN.

4.5. Some Related Specifications Under Development

The ATIS IPV Interoperability Forum (IIF), ATIS Cloud Services Forum (CSF), and the IETF CDNI Working Group are actively working on CDNI. The purpose of the IETF CDNI working group is to specify CDNI interfaces with associated protocols for control, logging, request routing and meta-data. As the IETF work progresses, the ATIS IIF and CSF architecture groups will be evaluating the protocols defined by the CDNI working group for inclusion in their respective CDN interconnect scenarios: IIF is focused on Internet Sourced Content and off net Delivery cases for IPTV, while CSF is focused on general Internet content distribution over cooperating CDN providers (ultimately leading to a multi-provider CDN federation).

[ATIS-0200003] describes the ATIS's CDN Use Case Specification and High Level Requirements. The current version describes the use case of off-peak software update download. Future revision will describe use cases of peak-hour file download and VOD streaming. The document defines two life-cycles: one for Primary & Supporting CDN, and the other for Content Provider, Primary & Supporting CDN.

The Primary & Supporting CDN life-cycle defines the requirements and steps for two CDN providers to enter a CDN peering agreement. It divides the life-cycle into four stages. The first stage is called "On Boarding". Two CDNs agree to interconnect. They will exchange capabilities and certification. The second stage is called "Active/Interconnected Environment". Two CDNs select and test the delivery functions and interfaces, SLA, and trouble-shooting procedures. The third and fourth stages are called "Termination Functions". Two CDNs agree to terminate the interconnection and implement the post mortem process.

Content Provider, Primary & Supporting CDN life-cycle defines the requirements and steps to establish the business relationship between Content Provider, Primary & Supporting CDN. It divides the life-cycle in five stages. The first and second stages are pre-sales and post-sales preparation for CDNI. It involves testing the delivery functions and exercising the infrastructure for CDNI services. The third stage is "Active/Interconnected Environment". The Content Provider will start delivering content to the Primary CDN, and the Primary CDN will use the Supporting CDN for user requests. In the fourth and fifth stages the interconnection will be terminated by one of the parties.

The IETF CDNI use cases specification [4] describes three sets of CDNI use cases: (1) Footprint Extension Use Cases, (2) Offload Use Cases and (3) CDN Capability Use Cases. The Footprint Extension Use Cases describe scenarios where a CDN provider wants to extend its reach to other region without compromising service quality. Offload Use Cases describe scenarios where a CDN provider uses CDNI to support overloading and resiliency. CDN Capability Use Cases describe scenarios where a CDN provider would use CDNI to improve quality of service.

The IETF CDNI requirements specification [5] describes the CDNI interface requirements. It contains five sets of requirements. The first set is Generic CDNI requirements. It defines the generic CDNI assumptions. For example: CDNI should be transparent to users. CDNI does not require the CDN providers to expose the internal topology information. The second set is CDNI Control Interface requirements. It defines the mission control requirements for two CDNs to form a CDNI. For example: the uCDN can mark or delete an object in the dCDN at any given time. The uCDN can also signal the dCDN to pre-position CDNI metadata. The third set is CDNI Request Routing Interface requirements. It defines the requirements to allow the uCDN to delegate the request to the dCDN. For example: uCDN must include necessary information such as user's origination and content identifier. The dCDN must signal the uCDN whether the request was completed or not; if not, what was the result. The fourth set is CDN Metadata Distribution Interface requirements. It defines the CDNI metadata delivery mechanism and necessary information. For example: the uCDN must indicate the location of the CDNI metadata to the dCDN. The CDNI metadata must contain enough information for the dCDN to make the delivery decision. The fifth set is CDNI Logging Interface requirements. It defines the requirements to allow two CDN providers to exchange logging information. This is particularly important for billing and trouble-shooting. For example: the chosen logging mechanism must be reliable and secure. The log file format must be common and simple.

The IETF description of the CDNI framework [6] provides examples of how CDNI should behave. It describes and lists the example of the two most common request redirections available today: DNS Redirection and HTTP Redirection.

Finally, the European Telecommunications Standards Institute) ETSI is focused on standardization from a system view (end to end). Both the ETSI MCD and TISPAN groups are involved in CDN. The TISPAN (Telecommunications and Internet converged Services and Protocols for Advanced Networking) technical committee is working on the CDN architecture and associated protocol adaption. The MCD (Media Content Distribution) technical committee is working on CDNI use-cases and requirements.

5. Conclusion

Recent Sandvine studies [7] show that OTT video distribution is becoming the most dominant source of traffic in the Internet. The growth of OTT video service in the next few years is predicted to be substantial. Content Providers often contract at most a few CDN Providers to stream the content via multiple ISPs to the end users for their consumption. Popular content can be delivered multiple times from the CDN provider to an ISP when it is requested by multiple users. Since an ISP only serves as an Internet provider, the ISP is unaware of the content being delivered, and cannot use their internal CDN to optimize the delivery.

CDNI is generating a lot of interest in the industry, since it allows CDNs from various parties to interconnect to each other. This enables ISP to cache the content to improve user experience and reduce packet duplication in the network. This document reviews the business motivations, use cases, high-level architecture and introduction to specification efforts being pursued by different standard bodies, to make CDNI a future reality.

6. References

[1] M. Day, B. Cain, G. Tomlinson and P. Rzewski, "A Model for Content Internetworking (CDI)", Internet Engineering Task Force RFC 3466, February 2003. www.ietf.org/rfc/rfc3466.txt

[2] Cooper, I., Melve, I. and G. Tomlinson, "Internet Web Replication and Caching Taxonomy", RFC 3040, June 2000. http://www.ietf.org/rfc/rfc3040.txt

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

[4] 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.

[5] Leung, K., Lee, Y., Le Faucheur, F., Viveganandhan, M., and G. Watson, "Content Distribution Network Interconnection (CDNI) Requirements", draft-lefaucheur-cdni-requirements-02 (work in progress), July 2011.

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

[7] Sandvine Spring 2011 Global Internet Phenomena Report, http://www.sandvine.com/downloads/documents/05-17-2011_phenomena/Sandvine%20Global%20Internet%20Phenomena%20Report.pdf

[8] Sandvine 2011 Global Internet Phenomena Spotlight – Netflix Rising, http://www.sandvine.com/downloads/documents/05-17-2011_phenomena/Sandvine%20Global%20Internet%20Phenomena%20Spotlight%20-%20Netflix%20Rising.pdf

[9] Jeff Baumgartner, Comcast's 'Project Infinity' Takes Flight, Light Reading Cable, http://www.lightreading.com/document.asp?doc_id=194827&site=lr_cable

[10] ATIS-0200003, "CDN Interconnection Use Case Specification and High Level Requirements", June 2011.