

Network Working Group
Internet-Draft
Intended status: Informational
Expires: July 12, 2013

R. van Brandenburg
O. van Deventer
TNO
F. Le Faucheur
K. Leung
Cisco Systems
January 8, 2013

**Models for adaptive-streaming-aware CDN Interconnection
draft-brandenburg-cdni-has-04**

Abstract

This document presents thoughts on the potential impact of supporting HTTP Adaptive Streaming technologies in CDN Interconnection (CDNI) scenarios. The intent is to present the authors' analysis of the CDNI-HAS problem space and discuss different options put forward both by the authors (and by others during informal discussions) on how to deal with HAS in the context of CDNI. This document has been used as input information during the WG process for making its decision regarding support for HAS.

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 12, 2013.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of

publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

| | | |
|--------------------------|---|--------------------|
| 1. | Introduction | 4 |
| 1.1. | Terminology | 4 |
| 2. | HTTP Adaptive Streaming aspects relevant to CDNI | 5 |
| 2.1. | Segmentation versus Fragmentation | 6 |
| 2.2. | Addressing chunks | 7 |
| 2.2.1. | Relative URLs | 8 |
| 2.2.2. | Absolute URLs with Redirection | 9 |
| 2.2.3. | Absolute URL without Redirection | 10 |
| 2.3. | Live vs. VoD | 10 |
| 2.4. | Stream splicing | 11 |
| 3. | Possible HAS Optimizations | 11 |
| 3.1. | File Management and Content Collections | 12 |
| 3.1.1. | General Remarks | 12 |
| 3.1.2. | Candidate approaches | 13 |
| 3.1.2.1. | Option 1.1: No HAS awareness | 13 |
| 3.1.2.2. | Option 1.2: Allow single file storage of fragmented content | 13 |
| 3.1.2.3. | Option 1.3: Access correlation hint | 14 |
| 3.1.3. | Recommendation | 15 |
| 3.2. | Content Acquisition of Content Collections | 15 |
| 3.2.1. | General Remarks | 15 |
| 3.2.2. | Candidate Approaches | 15 |
| 3.2.2.1. | Option 2.1: No HAS awareness | 15 |
| 3.2.2.2. | Option 2.2: Allow single file acquisition of fragmented content | 16 |
| 3.2.3. | Recommendation | 17 |
| 3.3. | Request Routing of HAS content | 17 |
| 3.3.1. | General remarks | 17 |
| 3.3.2. | Candidate approaches | 17 |
| 3.3.2.1. | Option 3.1: No HAS awareness | 17 |
| 3.3.2.2. | Option 3.2: Manifest File rewriting by uCDN | 19 |
| 3.3.2.3. | Option 3.3: Two-step Manifest File rewriting | 21 |
| 3.3.3. | Recommendation | 22 |
| 3.4. | Logging | 22 |
| 3.4.1. | General remarks | 22 |
| 3.4.2. | Candidate Approaches | 23 |
| 3.4.2.1. | Option 4.1: "Do-Nothing" Approach | 23 |
| 3.4.2.2. | Option 4.2: "CDNI Metadata Content Collection | |

| | |
|--|--------------------|
| ID" Approach | 25 |
| 3.4.2.3. Option 4.3: "CDNI Logging Interface Compression" Approach | 26 |
| 3.4.2.4. Option 4.4: "Full HAS awareness/per-Session-Logs" Approach | 27 |
| 3.4.3. Recommendation | 28 |
| 3.5. URL Signing | 30 |
| 3.5.1. HAS Implications | 31 |
| 3.5.2. CDNI Considerations | 31 |
| 3.5.3. Option 5.1: Do Nothing | 32 |
| 3.5.4. Option 5.2: Flexible URL Signing by CSP | 33 |
| 3.5.5. Option 5.3: Flexible URL Signing by Upstream CDN | 35 |
| 3.5.6. Option 5.4: Authorization Group ID and HTTP Cookie | 36 |
| 3.5.7. Option 5.5: HAS-awareness with HTTP Cookie in CDN | 37 |
| 3.5.8. Option 5.6: HAS-awareness with Manifest in CDN | 38 |
| 3.5.9. Recommendation | 39 |
| 3.6. Content Purge | 40 |
| 3.6.1. Option 6.1: No HAS awareness | 40 |
| 3.6.2. Option 6.2: Purge Identifiers | 40 |
| 3.6.3. Recommendation | 41 |
| 3.7. Other issues | 41 |
| 4. IANA Considerations | 42 |
| 5. Security Considerations | 42 |
| 6. Acknowledgements | 42 |
| 7. References | 42 |
| 7.1. Normative References | 42 |
| 7.2. Informative References | 42 |
| Authors' Addresses | 43 |

1. Introduction

[RFC6707] defines the problem space for CDN Interconnection (CDNI) and the associated CDNI interfaces. This includes support, through interconnected CDNs, of content delivery to endusers using HTTP progressive download and HTTP Adaptive Streaming (HAS).

HTTP Adaptive Streaming is an umbrella term for various HTTP-based streaming technologies that allow a client to adaptively switch between multiple bitrates depending on current network conditions. A defining aspect of HAS is that, since it is based on HTTP, it is a pull-based mechanism, with a client actively requesting content segments, instead of the content being pushed to the client by a server. Due to this pull-based nature, media servers delivering content using HAS often show different characteristics when compared with media servers delivering content using traditional streaming methods such as RTP/RTSP, RTMP and MMS.

This document presents a discussion on what the impact of these different characteristics is to the CDNI interfaces and what HAS-specific optimizations may be required or may be desirable. The scope of this document is to present the authors' analysis of the CDNI-HAS problem space and discuss different options put forward both by the authors (and by others during informal discussions) on how to deal with HAS in the context of CDNI. The document concludes by presenting the authors' recommendations on how the CDNI WG should deal with HAS in its initial charter, with a focus on 'making it work' instead of including 'nice-to-have' optimizations that might delay the development of the CDNI WG deliverables identified in its initial charter.

It should be noted that the document is not a WG document, but has been used as input information during the WG process for making its decision regarding support for HAS. We expect the analysis presented in the document will also be useful in the future if and when the WG re-charters and wants to re-assess the level of HAS optimizations to be supported in CDNI scenarios.

1.1. Terminology

This document uses the terminology defined in [RFC6707] and [I-D.ietf-cdni-framework].

For convenience, the definition of HAS-related terms are restated here:

Content Item: A uniquely addressable content element in a CDN. A content item is defined by the fact that it has its own Content

Metadata associated with it. It is the object of a request routing operation in a CDN. An example of a Content Item is a video file/stream, an audio file/stream or an image file.

Chunk: a fixed length element that is the result of a segmentation or fragmentation operation and that is independently addressable.

Fragment: A specific form of chunk (see [Section 2.1](#)). A fragment is stored as part of a larger file that includes all chunks that are part of the Chunk Collection.

Segment: A specific form of chunk (see [Section 2.1](#)). A segment is stored as a single file from a file system perspective.

Original Content: Non-chunked content that is the basis for a segmentation or fragmentation operation. Based on Original Content, multiple alternative representations (using different encoding methods, supporting different resolutions and/or targeting different bitrates) may be derived, each of which may be fragmented or segmented.

Chunk Collection: The set of all chunks that are the result of a single segmentation or fragmentation operation being performed on a single representation of the Original Content. A Chunk Collection is described in a Manifest File.

Content Collection: The set of all Chunk Collections that are derived from the same Original Content. A Content Collection may consist of multiple Chunk Collections, each corresponding to a single representation of the Original Content. A Content Collection may be described by one or more Manifest Files.

Manifest File: A Manifest File, also referred to as Media Presentation Description (MPD) file, is a file that list the way the content has been chunked (possibly for multiple encodings) and where the various chunks are located (in the case of segments) or how they can be addressed (in the case of fragments).

2. HTTP Adaptive Streaming aspects relevant to CDNI

In the last couple of years, a wide variety of HAS-like protocols have emerged. Among them are proprietary solutions such as Apple's HTTP Live Streaming (HLS), Microsoft's HTTP Smooth Streaming (HSS) and Adobe's HTTP Dynamic Streaming (HDS), and various standardized solutions such as 3GPP Adaptive HTTP Streaming (AHS) and MPEG Dynamic Adaptive Streaming over HTTP (DASH). While all of these technologies share a common set of features, each has its own defining elements.

This chapter will look at some of the common characteristics and some of the differences between these technologies and how those might be relevant to CDNI. In particular, [Section 2.1](#) will describe the various methods to store HAS content and [Section 2.2](#) will list three methods that are used to address HAS content in a CDN. After these generic HAS aspects are discussed, two special situations that need to be taken into account when discussing HAS are addressed: [Section 2.3](#) discusses the differences between Live and VoD content, while [Section 2.4](#) discusses the scenario where multiple streams are combined in a single Manifest File (e.g. for ad insertion purposes).

[2.1](#). Segmentation versus Fragmentation

All HAS implementations are based around a concept referred to as chunking: the concept of having a server split content up in numerous fixed duration chunks, which are independently decodable. By sequentially requesting and receiving chunks, a client can recreate and play out the content. An advantage of this mechanism is that it allows a client to seamlessly switch between different encodings of the same Original Content at chunk boundaries. Before requesting a particular chunk, a client can choose between multiple alternative encodings of the same chunk, irrespective of the encoding of the chunks it has requested earlier.

While every HAS implementation uses some form of chunking, not all implementations store the resulting chunks in the same way. In general, there are two distinct methods of performing chunking and storing the results: segmentation and fragmentation.

- With segmentation, which is for example mandatory in all versions of Apple's HLS prior to version 7, the chunks, in this case also referred to as segments, are stored completely independent from each other, with each segment being stored as a separate file from a file system perspective. This means that each segment has its own unique URL with which it can be retrieved.
- With fragmentation (or virtual segmentation), which is for example used in Microsoft's Smooth Streaming, all chunks, or fragments, belonging to the same Chunk Collection are stored together, as part of a single file. While there are a number of container formats which allow for storing this type of chunked content, Fragmented MP4 is most commonly used. With fragmentation, a specific chunk is addressable by subfixing the common file URL with an identifier uniquely identifying the chunk that one is interested in, either by timestamp, by byterange, or in some other way.

While one can argue about the merits of each of these two different

methods of handling chunks, both have their advantages and drawbacks in a CDN environment. For example, fragmentation is often regarded as a method that introduces less overhead, both from a storage and processing perspective. Segmentation on the other hand, is regarded as being more flexible and easier to cache. In practice, current HAS implementations increasingly support both methods.

[2.2.](#) Addressing chunks

In order for a client to request chunks, either in the form of segments or in the form of fragments, it needs to know how the content has been chunked and where to find the chunks. For this purpose, most HAS protocols use a concept that is often referred to as a Manifest File (also known as Media Presentation Description, or MPD); i.e. a file that lists the way the content has been chunked and where the various chunks are located (in the case of segments) or how they can be addressed (in the case of fragments). A Manifest File, or set of Manifest Files, may also identify the different representations, and thus Chunk Collections, available for a content.

In general, a HAS client will first request and receive a Manifest File, and then, after parsing the information in the Manifest File, proceed with sequentially requesting the chunks listed in the Manifest File. Each HAS implementation has its own Manifest File format and even within a particular format there are different methods available to specify the location of a chunk.

Of course managing the location of files is a core aspect of every CDN, and each CDN will have its own method of doing so. Some CDNs may be purely cache-based, with no higher-level knowledge of where each file resides at each instant in time. Other CDNs may have dedicated management nodes which, at each instant in time, do know at which servers each file resides. The CDNI interfaces designed in the CDNI WG will probably need to be agnostic to these kinds of CDN-internal architecture decisions. In the case of HAS there is a strict relationship between the location of the content in the CDN (in this case chunks) and the content itself (the locations specified in the Manifest File). It is therefore useful to have an understanding of the different methods in use in CDNs today for specifying chunk locations in Manifest Files. The different methods for doing so are described in sections [2.2.1](#) to [2.2.3](#).

Although these sections are especially relevant for segmented content, due to its inherent distributed nature, the discussed methods are also applicable to fragmented content. Furthermore, it should be noted that the methods detailed below for specifying locations of content items in Manifest Files do not only relate to temporally segmented content (e.g. segments and fragments), but are

also relevant in situations where content is made available in multiple representations (e.g., in different qualities, encoding methods, resolutions and/or bitrates). In this case the content consists of multiple chunk collections, which may be described by either a single Manifest File or multiple interrelated Manifest Files. In the latter case, there may be a high-level Manifest File describing the various available bitrates, with URLs pointing to separate Manifest Files describing the details of each specific bitrate. For specifying the locations of the other Manifest Files, the same methods apply that are used for specifying chunk locations.

One final note relates to the delivery of the Manifest Files themselves. While in most situations the delivery of both the Manifest File and the chunks are handled by the CDN, there are scenarios imaginable in which the Manifest File is delivered by e.g. the Content Provider, and the Manifest File is therefore not visible to the CDN.

2.2.1. Relative URLs

One method for specifying chunk locations in a Manifest File is through the use of relative URLs. A relative URL is a URL that does not include the HOST part of a URL but only includes (part of) the PATH part of a URL. In practice, a relative URL is used by the client as being relative to the location where the Manifest File has been acquired from. In these cases a relative URL will take the form of a string that has to be appended to the location of the Manifest File to get the location of a specific chunk. This means that in the case a Manifest File with relative URLs is used, all chunks will be delivered by the same surrogate that delivered the Manifest File. A relative URL will therefore not include a hostname.

For example, in the case a Manifest File has been requested (and received) from:

http://surrogate.server.cdn.example.com/content_1/manifest.xml

, a relative URL pointing to a specific segment referenced in the Manifest File might be:

segments/segment1_1.ts

Which means that the client should take the location of the Manifest File and append the relative URL. In this case, the segment would then be requested from http://surrogate.server.cdn.example.com/content_1/segments/segment1_1.ts

The downside of using relative URLs is that it forces a CDN to

deliver all segments belonging to a given content item with the same surrogate that delivered the Manifest File for that content item, which results in limited flexibility. Another drawback is that Relative URLs do not allow for fallback URLs; should the surrogate that delivered the Manifest File break down, the client is no longer able to request chunks. The advantage of relative URLs is that it is very easy to transfer content between different surrogates and even CDNs.

2.2.2. Absolute URLs with Redirection

Another method for specifying locations of chunks (or other Manifest Files) in a Manifest File is through the use of an absolute URL. An absolute URL contains a fully formed URL (i.e. the client does not have to calculate the URL as in the case of the relative URL but can use the URL from the Manifest File directly).

In the context of Manifest Files, there are two types of absolute URLs imaginable: Absolute URLs with Redirection and Absolute URLs without Redirection. The two methods differ in whether the URL points to a request routing node which will redirect the client to a surrogate (Absolute URL with Redirection) or point directly to a surrogate hosting the requested content (Absolute URL without Redirection).

In the case of Absolute URLs with Redirection, a request for a chunk is handled by the request routing system of a CDN just as if it were a standalone (non-HAS) content request, which might include looking up the surrogate (and/or CDN) best suited for delivering the requested chunk to the particular user and sending an HTTP redirect to the user with the URL pointing to the requested chunk on the specified surrogate (and/or CDN), or a DNS response pointing to the specific surrogate.

An example of an Absolute URL with Redirection might look as follows:

[http://requestrouting.cdn.example.com/
content_request?content=content_1&segment=segment1_1.ts](http://requestrouting.cdn.example.com/content_request?content=content_1&segment=segment1_1.ts)

As can be seen from this example URL, the URL includes a pointer to a general CDN request routing function and includes some arguments identifying the requested segment.

The advantage of using Absolute URLs with Redirection is that it allows for maximum flexibility (since chunks can be distributed across surrogates and CDN in any imaginable way) without having to modify the Manifest File every time one or more chunks are moved (as is the case when Absolute URLs without Redirection are used). The

downside of this method is that it can add significant load to a CDN request routing system, since it has to perform a redirect every time a client requests a new chunk.

2.2.3. Absolute URL without Redirection

In the case of the Absolute URL without Redirection, the URL points directly to the specific chunk on the actual surrogate that will deliver the requested chunk to the client. In other words, there will be no HTTP redirection operation taking place between the client requesting the chunk and the chunk being delivered to the client by the surrogate.

An example of an Absolute URL without Redirection is the following:

http://surrogate.cdn.example.com/content_1/segments/segment1_1.ts

As can be seen from this example URL, the URL includes both the identifier of the requested segment (in this case segment1_1.ts), as well as the server that is expected to deliver the segment (in this case surrogate.cdn.example.com). With this, the client has enough information to directly request the specific segment from the specified surrogate.

The advantage of using Absolute URLs without Redirection is that it allows more flexibility compared to using Relative URLs (since segments do not necessarily have to be delivered by the same server) while not requiring per-segment redirection (which would add significant load to the node doing the redirection). The drawback of Absolute URLs without Redirection is that it requires a modification of the Manifest File every time content is moved to a different location (either within a CDN or across CDNs).

2.3. Live vs. VoD

Though the formats and addresses of Manifest Files and chunk files do not typically differ significantly between live and Video-on-Demand (VoD) content, the time at which the Manifest Files and chunk files become available does differ significantly. For live content, chunk files and their corresponding Manifest Files are created and delivered in real-time. This poses a number of potential issues for HAS optimization:

- With live content, chunk files are made available in real-time. This limits the applicability of bundling for content acquisition purposes. Prepositioning may still be employed, however, any significant latency in the prepositioning may diminish the value of prepositioning if a client requests the chunk prior to

prepositioning, or if the prepositioning request is serviced after the chunk playout time has passed.

- In the case of live content, Manifest Files must be updated for each chunk and therefore must be retrieved by the client prior to each chunk request. Any Manifest-File based optimization schemes must therefore be prepared to optimize on a per-segment request basis. Manifest Files may also be polled multiple times prior to the actual availability of the next chunk.
- Since live Manifest Files are updated as each new chunk becomes available, the cacheability of Manifest Files is limited. Though timestamping and reasonable TTLs can improve delivery performance, timely replication and delivery of updated Manifest Files is critical to ensuring uninterrupted playback.
- Manifest Files are typically updated after the corresponding chunk is available for delivery, to prevent premature requests for chunks which are not yet available. HAS optimization approaches which employ dynamic Manifest File generation must be synchronized with chunk creation to prevent playback errors.

2.4. Stream splicing

Stream splicing is used to create media mashups, combining content from multiple sources. A common example in which content resides outside the CDNs is with advertisement insertion, for both VoD and live streams. Manifest Files which contain Absolute URLs with redirection may contain chunk or nested Manifest File URLs which point to content not delivered via any of the interconnected CDNs.

Furthermore, client and downstream proxy devices may depend on non-URL information provided in the Manifest File (e.g., comments or custom tags) for performing stream splicing. This often occurs outside the scope of the interconnected CDNs. HAS optimization schemes which employ dynamic Manifest File generation or rewriting must be cognizant of chunk URLs, nested Manifest File URLs, and other metadata which should not be modified or removed. Improper modification of these URLs or other metadata may cause playback interruptions, and in the case of unplayed advertisements, may result in loss of revenue for content providers.

3. Possible HAS Optimizations

In the previous chapter, some of the unique properties of HAS have been discussed. Furthermore, some of the CDN-specific design decisions with regards to addressing chunks have been detailed. In

this chapter, the impact of supporting HAS in CDN Interconnection scenarios will be discussed.

There are a number of topics, or problem areas, that are of particular interest when considering the combination of HAS and CDNI. For each of these problem areas it holds that there are a number of different ways in which the CDNI Interfaces can deal with them. In general it can be said that each problem area can either be solved in a way that minimizes the amount of HAS-specific changes to the CDNI Interfaces or in way that maximizes the flexibility and efficiency with which the CDNI Interfaces can deliver HAS content. The goal for the CDNI WG should probably be to try to find the middle ground between these two extremes and try to come up with solutions that optimize the balance between efficiency and additional complexity.

In order to allow the WG to make this decision, this chapter will briefly describe each of the following problem areas together with a number of different options for dealing with them. [Section 3.1](#) will discuss the problem of how to deal with file management of groups of files, or Content Collections. [Section 3.2](#) will deal with a related topic: how to do content acquisition of Content Collections between the uCDN and dCDN. After that, [Section 3.3](#) describes the various options for the request routing of HAS content, particularly related to Manifest Files. [Section 3.4](#) talks about a number of possible optimizations for the logging of HAS content, while [Section 3.5](#) discusses the options regarding URL signing. [Section 3.6](#) finally, describes different scenarios for dealing with the removal of HAS content from CDNs.

[3.1. File Management and Content Collections](#)

[3.1.1. General Remarks](#)

One of the unique properties of HAS content is that it does not consist of a single file or stream but of multiple interrelated files (segment, fragments and/or Manifest Files). In this document this group of files is also referred to as a Content Collection. Another important aspect is the difference between segments and fragments (see [Section 2.1](#)).

Irrespective of whether segments or fragments are used, different CDNs might handle Content Collections differently from a file management perspective. For example, some CDNs might handle all files belonging to a Content Collection as individual files, which are stored independently from each other. An advantage of this approach is that makes it easy to cache individual chunks. Other CDNs might store all fragments belonging to a Content Collection in a bundle, as if they were a single file (e.g. by using a fragmented MP4

container). The advantage of this approach is that it reduces file management overhead.

This section will look at the various ways with which the CDNI interfaces might deal with these differences in handling Content Collections from a file management perspective. The different options can be distinguished based on the level of HAS-awareness they require on the part of the different CDNs and the CDNI interfaces.

3.1.2. Candidate approaches

3.1.2.1. Option 1.1: No HAS awareness

This first option assumes no HAS awareness in both the involved CDNs and the CDNI Interfaces. This means that the uCDN uses individual files and the dCDN is not explicitly made aware of the relationship between chunks and it doesn't know which files are part of the same Content Collection. In practice this scenario would mean that the file management method used by the uCDN is simply imposed on the dCDN as well.

This scenario also means that it is not possible for the dCDN to use any form of file bundling, such as the single-file mechanism which can be to store fragmented content as a single file (see [Section 2.1](#)). The one exception to this rule is the situation where the content is fragmented and the Manifest Files on the uCDN contains byte range requests, in which case the dCDN might be able to acquire fragmented content as a single file (see [Section 3.2.2.2](#)).

Effect on CDNI interfaces:

- o None

Advantages/Drawbacks:

- + No HAS awareness necessary in CDNs, no changes to CDNI Interfaces necessary
- The dCDN is forced to store chunks as individual files.

3.1.2.2. Option 1.2: Allow single file storage of fragmented content

In some cases, the dCDN might prefer to store fragmented content as a single file on its surrogates to reduce file management overhead. In order to do so, it needs to be able to either acquire the content as a single file (see [Section 3.2.2.2](#)), or merge the different chunks together and place them in the same container (e.g. fragmented MP4). The downside of this is that in order to do so, the dCDN needs to be

fully HAS aware.

Effect on CDNI interfaces:

- o CDNI Metadata Interface: Add fields for indicating the particular type of HAS (e.g. MPEG DASH or HLS) that is used and whether segments or fragments are used
- o CDNI Metadata Interface: Add field for indicating the name and type of the Manifest File(s)

Advantages/Drawbacks:

- + Allows dCDN to store fragmented content as a single file, reducing file management overhead
- Complex operation, requiring dCDN to be fully HAS aware

3.1.2.3. Option 1.3: Access correlation hint

An intermediary approach between the two extremes detailed in the previous two sections is one that uses a 'Access Correlation Hint'. This hint, which is added to the CDNI Metadata of all chunks of a particular Content Collection, indicates that those files are likely to be requested in a short time window from each other. This information can help a dCDN to implement local file storage optimizations for VoD items (e.g. by bundling all files with the same Access Correlation Hint value in a single bundle/file), thereby reducing the number of files it has to manage while not requiring any HAS awareness.

Effect on CDNI interfaces:

- o CDNI Metadata Interface: Add field for indicating Access Correlation Hint

Advantages/Drawbacks:

- + Allows dCDN to perform file management optimization
- + Does not require any HAS awareness
- + Very small impact on CDNI Interfaces
- Expected benefit compared with Option 1.1 is small

3.1.3. Recommendation

Based on the listed pros and cons, the authors recommend the WG to go for Option 1.1, the 'Do Nothing'-approach. The likely benefits from going for Option 1.3 are not believed to be significant enough to warrant changing the CDNI Metadata Interface. Although Option 1.2 would bring definite benefits for HAS aware dCDNs, going for this options would require significant CDNI extensions that would impact the WG's milestones. The authors therefore don't recommend to include it in the current work but mark it as a possible candidate for rechartering once the initial CDNI solution is completed.

3.2. Content Acquisition of Content Collections

3.2.1. General Remarks

In the previous section the relationship between file management and HAS in a CDNI scenario has been discussed. This section will discuss a related topic, which is content acquisition between two CDNs.

With regards to content acquisition, it is important to note the difference between CDNs that do Dynamic Acquisition of content and CDNs that perform Content Pre-positioning. In the case of dynamic acquisition, a CDN only requests a particular content item when a cache-miss occurs. In the case of pre-positioning, a CDN pro-actively places content items on the nodes on which it expects traffic for that particular content item. For each of these types of CDNs, there might be a benefit in being HAS aware. For example, in the case of dynamic acquisition, being HAS aware means that after a cache miss for a giving chunk occurs, that node might not only acquire the requested chunk, but might also acquire some related chunks that are expected to be requested in the near future. In the case of pre-positioning, similar benefits can be had.

3.2.2. Candidate Approaches

3.2.2.1. Option 2.1: No HAS awareness

This first option assumes no HAS awareness in both the involved CDNs and the CDNI Interfaces. Just as with Option 1.1 discussed in the previous section with regards to file management, having no HAS awareness means that the dCDN is not aware of the relationship between chunks. In the case of content acquisition, this means that each and every file belonging to a Content Collection will have to be individually acquired from the uCDN by the dCDN. The exception to the rule is in cases with fragmented content where the uCDN uses Manifest Files which contain byte range requests. In this case the dCDN can simply omit the byte range identifier and acquire the

complete file.

The advantage of this approach is that it is highly flexible. If a client only requests a small portion of the chunks belonging to a particular Content Collection, the dCDN only has to acquire those chunks from the uCDN, saving both bandwidth and storage capacity.

The downside of acquiring content on a per-chunk basis is that it creates more transaction overhead between the dCDN and uCDN compared to a method in which entire Content Collections can be acquired as part of one transaction.

Effect on CDNI interfaces:

- o None

Advantages/Drawbacks:

- + Per-chunk content acquisition allows for high level of flexibility between dCDN and uCDN
- Per-chunk content acquisition creates more transaction overhead between dCDN and uCDN

3.2.2.2. Option 2.2: Allow single file acquisition of fragmented content

As discussed in [Section 3.2.2.1](#), there is one (fairly rare) case where fragmented content can be acquired as a single file without any HAS awareness and that is when fragmented content is used and where the Manifest File specifies byte range request. This section discusses how to perform single file acquisition in the other (very common) cases. To do so, the dCDN would have to have full-HAS awareness (at least to the extent of being able to map between single file and individual chunks to serve).

Effect on CDNI interfaces:

- o CDNI Metadata Interface: Add fields for indicating the particular type of HAS (e.g. MPEG DASH or HLS) that is used and whether segments or fragments are used
- o CDNI Metadata Interface: Add field for indicating the name and type of the Manifest File(s)

Advantages/Drawbacks:

- + Allows for more efficient content acquisition in all HAS-specific supported forms
- Requires full HAS awareness on part of dCDN
- Requires significant CDNI Metadata Interface extensions

3.2.3. Recommendation

Based on the listed pros and cons, the authors recommend the WG to go for Option 2.1 since it is sufficient to 'make HAS work'. While Option 2.2 would bring benefits to the acquisition of large Content Collections, it would require significant CDNI extensions which would impact the WG's milestones. Option 2.2 might be a candidate to include in possible rechartering once the initial CDNI solution is completed.

3.3. Request Routing of HAS content

3.3.1. General remarks

In this section the effect HAS content has on request routing will be identified. Of particular interest in this case are the different types of Manifest Files that might be used. In [Section 2.2](#), three different methods for identifying and addressing chunks from within a Manifest File were described: Relative URLs, Absolute URLs without Redirection and Absolute URLs with Redirection. Of course not every current CDN will use and/or support all three methods. Some CDNs may only use one of the three methods, while others may support two or all three.

An important factor in deciding which chunk addressing method is used is the Content Provider. Some Content Providers may have a strong preference for a particular method and deliver the Manifest Files to the CDN in a particular way. Depending on the CDN and the agreement it has with the Content Provider, a CDN may either host the Manifest Files as they were created by the Content Provider, or modify the Manifest File to adapt it to its particular architecture (e.g. by changing relative URLs to Absolute URLs which point to the CDN Request Routing function).

3.3.2. Candidate approaches

3.3.2.1. Option 3.1: No HAS awareness

This first option assumes no HAS awareness in both the involved CDNs and the CDNI Interfaces. This scenario also assumes that neither the dCDN nor the uCDN have the ability to actively manipulate Manifest

Files. As was also discussed with regards to file management and content acquisition, having no HAS awareness means that each file constituting a Content Collections is handled on an individual basis, with the dCDN unaware of any relationship between files.

The only chunk addressing method that works without question in this case is Absolute URLs with Redirection. In other words, the Content Provider that ingested the content into the uCDN created a Manifest File with each chunk location pointing to the Request Routing function of the uCDN. Alternatively, the Content Provider may have ingested the Manifest File containing relative URLs and the uCDN ingestion function has translated these to Absolute URLs pointing to the Request Routing function.

In this Absolute URL with Redirection case, the uCDN can simply have the Manifest File be delivered by the dCDN as if it were a regular file. Once the client parses the Manifest File, it will request any subsequent chunks from the uCDN Request Routing function. That function can then decide to outsource the delivery of that chunk to the dCDN. Depending on whether HTTP-based (recursive or iterative) or DNS-based request routing is used, the uCDN Request Routing function will then either directly or indirectly redirect the client to the Request Routing function of the dCDN (assuming it does not have the necessary information to redirect the client directly to a surrogate in the dCDN).

The drawback of this method is that it creates a large amount of request routing overhead for both the uCDN and dCDN. For each chunk the full inter-CDN Request Routing process is invoked (which can result in two HTTP redirections in the case of iterative redirection, or result in one HTTP redirection plus one CDNI Request Routing/Redirection Interface request/response). Even in the case where DNS-based redirection is used, there might be significant overhead involved since both the dCDN and uCDN Request Routing function might have to perform database lookups and query each other. While with DNS this overhead might be reduced by using DNS' inherent caching mechanism, this will have significant impact on the accuracy of the redirect.

With no HAS awareness, Relative URLs might or might not work depending on the type of Relative URL that is used. When a uCDN delegates the delivery of a Manifest File containing Relative URLs to a dCDN, the client goes directly to the dCDN surrogate from which it has received the Manifest File for every subsequent chunk. As long as the Relative URL is not path-absolute (see [[RFC3986](#)]), this approach will work fine.

Since using Absolute URLs without Redirection inherently require a

HAS aware CDN, they also cannot be used in this case. The reason for this is that with Absolute URLs without Redirection, the URLs in the Manifest File will point directly to a surrogate in the uCDN. Since this scenario assumes no HAS awareness on the part of the dCDN or uCDN, it is impossible for either of these CDNs to rewrite the Manifest File and thus allow the client to either go to a surrogate in the dCDN or to a request routing function.

Effect on CDNI interfaces:

- o None

Advantages/Drawbacks:

- + Supports Absolute URLs with Redirection
- + Supports Relative URLs
- + Does not require HAS awareness and/or changes to the CDNI Interfaces
- Not possible to use Absolute URLs without Redirection
- Creates significant signaling overhead in case Absolute URLs with Redirection are used (inter-CDN request redirection for each chunk)

3.3.2.2. Option 3.2: Manifest File rewriting by uCDN

While Option 3.1 does allow for Absolute URLs with Redirection to be used, it does so in a way that creates a high-level of request routing overhead for both the dCDN and the uCDN. This option presents a solution to significantly reduce this overhead.

In this scenario, the uCDN is able to rewrite the Manifest File (or generate a new one) to be able to remove itself from the request routing chain for chunks being referenced in the Manifest File. As described in [Section 3.3.2.1](#), in the case of no HAS awareness the client will go to the uCDN request routing function for each chunk request. This request routing function can then redirect the client to the dCDN request routing function. By rewriting the Manifest File (or generating a new one), the uCDN is able to remove this first step, and have the Manifest File point directly to the dCDN request routing function.

A key advantage of this solution is that it does not directly have an impact on the CDNI Interfaces and is therefore transparent to these interfaces. It is a CDN-internal function that a uCDN can perform

autonomously by using information configured for regular CDNI operation or that is received from the dCDN as part of the regular communication using the CDNI Request Routing/Redirection Interface.

More specifically, in order for the uCDN to rewrite the Manifest File, the minimum information needed is the location of the dCDN request routing function (or alternatively the location of the dCDN delivering surrogate). This information can be available from configuration or can be derived from the regular CDNI Request Routing/Redirection Interface. For example, the uCDN may ask the dCDN for the location of its request routing node (through the CDNI Request Routing/Redirection Interface) every time a request for a Manifest File is received and processed by the uCDN request routing function. The uCDN would then modify the Manifest File and deliver the Manifest File to the client. One advantage of this method is that it maximizes efficiency and flexibility by allowing the dCDN to optionally respond with the locations of its surrogates instead of the location of its request routing function (and effectively turning the URLs into Absolute URLs without Redirection). There are many variations around this approach, such as where the modification of the Manifest File is only performed once (or once per period of time) by the uCDN request routing function, when the first client for that particular Content Collection (and redirected to that particular dCDN) sends a Manifest File request. The advantage of such a variation is that the uCDN only has to modify the Manifest File once (or once per time period). The drawback of this variation is that the dCDN is no longer in a position to influence the request routing decision across individual content requests.

It should be noted that there are a number of things to take into account when changing a Manifest File (see for example [Section 2.3](#) and [Section 2.4](#) on live HAS content and ad insertion). Furthermore, some Content Providers might have issues with a CDN changing Manifest Files. However, in this option the Manifest File manipulation is only being performed by the uCDN, which can be expected to be aware of these limitations if it wants to perform Manifest File manipulation since it is in its own best interest that its customer's content gets delivered in the proper way and since there is a direct commercial and technical relationship between the uCDN (the Authoritative CDN in this scenario) and its customer (the Content Provider). Should the Content Provider want to limit Manifest File manipulation, it can simply arrange this with the uCDN bilaterally.

Effect on CDNI interfaces:

- o None

Advantages/Drawbacks:

- + Possible to significantly decrease signaling overhead when using Absolute URLs.
- + (Optional) Possible to have uCDN rewrite the Manifest File with locations of surrogates in dCDN (turning Absolute URLs with Redirection in Absolute URLs without Redirection)
- + No changes to CDNI Interfaces
- + Does not require HAS awareness in dCDN
- Requires high level of HAS awareness in uCDN (for modifying Manifest Files)

3.3.2.3. Option 3.3: Two-step Manifest File rewriting

One of the possibilities with Option 3.2 is allowing the dCDN to provide the locations of a specific surrogate to the uCDN, so that the uCDN can fit the Manifest File with Absolute URLs without Redirection and the client can request chunks directly from a dCDN surrogate. However, some dCDNs might not be willing to provide this information to the uCDN. In that case they can only provide the uCDN with the location of their request routing function and thereby preventing use of Absolute URLs without Redirection.

One method for solving this limitation is allowing two-step Manifest File manipulation. In the first step the uCDN would perform its own modification, and place the locations of the dCDN request routing function in the Manifest File. Then, once a request for the Manifest File comes in at the dCDN request routing function, it would perform a second modification in which it replaces the URLs in the Manifest Files with the URLs of its surrogates. This way the dCDN can still profit from having limited request routing traffic, while not having to share sensitive surrogate information with the uCDN.

The downside of this approach is that it not only assumes HAS awareness in the dCDN but it also requires some HAS-specific additions to the CDNI Metadata Interface. In order for the dCDN to be able to change the Manifest File, it has to have some information about the structure of the content. Specifically, it needs to have information about which chunks make up the Content Collection.

Effect on CDNI interfaces (apart from those already listed under Option 3.2):

- o CDNI Metadata Interface: Add necessary fields for conveying HAS specific information (e.g. the files that make up the Content Collection) to the dCDN.

- o CDNI Metadata Interface: Allow dCDN to modify Manifest File

Advantages/Drawbacks (apart from those already listed under Option 3.2):

- + Allows dCDN to use Absolute URLs without Redirection without having to convey sensitive information to the uCDN
- Requires high level of HAS awareness in dCDN (for modifying Manifest Files)
- Requires adding HAS-specific and Manifest File manipulation specific information to the CDNI Metadata Interface

3.3.3. Recommendation

Based on the listed pros and cons, the authors recommend to go for Option 3.1, with Option 3.2 as an optional feature that may be supported as a CDN-internal behavior by an uCDN. While Option 3.1 allows for HAS content to be delivered using the CDNI interfaces, it does so with some limitations regarding supported Manifest Files and, in some cases, with large signaling overhead. Option 3.2 can solve most of these limitations and presents a significant reduction of the request routing overhead. Since Option 3.2 does not require any changes to the CDNI interfaces but only changes the way the uCDN uses the existing interfaces, supporting it is not expected to result in a significant delay of the WG's milestones. The authors recommend the WG to not include Option 3.3, since it raises some questions of potential brittleness and including it would result in a significant delay of the WG's milestones.

3.4. Logging

3.4.1. General remarks

As stated in [[RFC6707](#)], "the CDNI Logging interface enables details of logs or events to be exchanged between interconnected CDNs".

As discussed in [I-D.[draft-bertrand-cdni-logging](#)], the CDNI logging information can be used for multiple purposes including maintenance/debugging by uCDN, accounting (e.g. in view of billing or settlement), reporting and management of end-user experience (e.g. to the CSP), analytics (e.g. by the CSP) and control of content distribution policy enforcement (e.g. by the CSP).

The key consideration for HAS with respect to logging is the potential increase of the number of Log records by two to three orders of magnitude, as compared to regular HTTP delivery of a video,

since, by default, log records would typically be generated on a per-chunk-delivery basis instead of per-content-item-delivery basis. This impacts the scale of every processing step in the Logging Process (see [I-D.[draft-bertrand-cdni-logging](#)]), including:

- a. Logging information generation and storing on CDN elements (Surrogate, Request Routers,..)
- b. Logging information aggregation within a CDN
- c. Logging information manipulation (including information protection, filtering , update and rectification)
- d. (Where needed) Logging information CDNI reformatting (e.g. reformatting from CDN-specific format to the CDNI Logging Interface format for export by dCDN to uCDN)
- e. Logging exchange via CDNI Logging Interface
- f. (Where needed) Logging re-reformatting (e.g. reformatting from CDNI Logging Interface format into log-consuming specific application)
- g. Logging consumption/processing (e.g. feed logs into uCDN accounting application, feed logs into uCDN reporting system to provide per CSP views, feed logs into debugging tools)

Note that there may be multiple instances of step [f] and [g] running in parallel.

While the CDNI Logging Interface is only used to perform step [e], we note that its format directly affects step [d] and [f] and that its format also indirectly affects the other steps (for example if the CDNI Logging Interface requires per-chunk log records, step [a], [b] and [d] cannot operate on a per-HAS-session basis and they also need to operate on a per-chunk basis).

This section discusses the main candidate approaches identified for CDNI in terms of dealing with HAS with respect to Logging.

3.4.2. Candidate Approaches

3.4.2.1. Option 4.1: "Do-Nothing" Approach

In this approach nothing is done specifically for HAS so that each HAS-chunk delivery is considered, for CDNI Logging, as a standalone content delivery. In particular, a separate log record for each HAS-chunk delivery is included in the CDNI Logging Interface in step [e]

(as defined in [Section 3.4.1](#)). This approach requires that step [a], [b], [c], [d] and [e] also be performed on a per-chunk basis. This approach allows [g] to be performed either on a per-chunk basis (assuming step [f] maintains per-chunk records) or on a more "summarized" manner such as per-HAS-Session basis (assuming step [f] summarizes per-chunk records into per-HAS-session records).

Effect on CDNI interfaces:

- o None

Advantages/Drawbacks:

- + No information loss (i.e. all details of each individual chunk delivery are preserved). While this full level of detail may not be needed for some Log consuming applications (e.g. billing), this full level of detail is likely valuable (possibly required) for some Log consuming applications (e.g. debugging)
- + Easier integration (at least in the short term) into existing Logging tools since those are all capable of handling per-chunk records
- + No extension needed on CDNI interfaces
- High volume of logging information to be handled (storing & processing) at every step of the Logging process from [a] to [g] (while summarization in step [f] is conceivable, it may be difficult to achieve in practice without any hints for correlation in the log records).

An interesting question is whether a dCDN could use the CDNI Logging interface specified for the "Do-Nothing" approach to report summarized "per-session" log information in the case where the dCDN performs such summarization. The high level idea would be that, when a dCDN performs HAS log summarization for its own purposes anyways, this dCDN could include, in the CDNI Logging interface, one (or a few) log entry for a HAS session (instead of one entry per HAS-chunk) that summarizes the deliveries of many/all HAS-chunk for a session. However, the authors feel that, when considering the details of this, this is not achievable without explicit agreement between the uCDN and dCDN about how to perform/interpret such summarization. For example, when a HAS session switches between representations, the uCDN and dCDN would have to agree on things such as:

- o whether the session will be represented by a single log entry (which therefore cannot convey the distribution across representations) or multiple log entries such as one entry per

contiguous period at a given representation (which therefore would be generally very difficult to correlate back into a single session)

- o what would the single URI included in the log entry correspond to? the Manifest/top-level-playlist/next-level-playlist,...

The authors feel that since explicit agreement is needed between uCDN and dCDN on how to perform/interpret the summarization, then, to this can only work if specified as part of the CDNI Logging interface and then effectively boils down to Option 4.4 defined below ("Full HAS awareness/per-Session-Logs" Approach).

We note that support by CDNI of a mechanism (independent of HAS) allowing the customization of the fields to be reported in log entries by the dCDN to the uCDN would have a mitigation effect on the HAS logging scaling concerns because it ensures that only the necessary subset of fields are actually stored, reported and processed.

3.4.2.2. Option 4.2: "CDNI Metadata Content Collection ID" Approach

In this approach, a "Content Collection IDentifier (CCID)" field is distributed through the CDNI Metadata Interface and the same CCID value is associated through the CDNI Metadata interface with every chunk of the same Content Collection. The CCID value needs to be such that it allows, in combination with the content URI, to uniquely identify a Content Collection. When distributed, and CCID logging is requested from the dCDN, the dCDN Surrogates are to store the CCID value in the corresponding log entries. The objective of this field is to facilitate optional summarization of per-chunk records at step [f] into something along the lines of per-HAS-session logs, at least for the Log consuming applications that do not require per-chunk detailed information (for example billing).

We note that, if the downstream CDN happens to have sufficient HAS awareness to be able to generate a "Session IDentifier (Session-ID)", optionally including such Session-ID (in addition to the CCID) in the per-chunk log record would further facilitate optional summarization performed at step [f]. The Session-ID value to be included in a log record by the delivering CDN is such that

- o different per-chunk log records with the same Session-ID value must correspond to the same user session (.i.e delivery of same content to same enduser at a given point in time).
- o log records for different chunks of the same user session (.i.e delivery of same content to same enduser at a given point in time)

should be provided with the same session-ID value. While undesirable, there may be situations where the delivering CDN uses more than one session-ID value for different per-chunk log records of a given session, for example in scenarios of fail-over or load-balancing across multiple Surrogates and where the delivering CDN does not implement mechanisms to synchronize session-IDs across Surrogates.

Effect on CDNI interfaces:

- o CDNI Metadata interface: One additional metadata field (CCID) in CDNI Metadata Interface. We note that a similar Content Collection ID is discussed for handling of other aspects of HAS and observe that further thought is needed to determine whether such CCID should be shared for multiple purposes or should be independent.
- o CDNI Logging interface: Two additional fields (CCID and Session-ID) in CDNI Logging records.

Advantages/Drawbacks:

- + No information loss (i.e. all details of each individual chunk delivery are preserved). While this full level of detail may not be needed for some Log consuming applications (e.g. billing), this full level of detail is likely valuable (possibly required) for some Log consuming applications (e.g. debugging)
- + Easier integration (at least in the short term) into existing Logging tools since those are all capable of handling per-chunk records
- + Very minor extension to CDNI interfaces needed
- + Facilitated summarization of records related to a HAS session in step [f] and therefore ability to operate on lower volume of logging information in step [g] by log consuming applications that do not need per-chunk record details (e.g. billing) or that need per-session information (e.g. analytics)
- High volume of logging information to be handled (storing & processing) at every step of the Logging process from [a] to [f].

3.4.2.3. Option 4.3: "CDNI Logging Interface Compression" Approach

In this approach, a loss-less compression technique is applied to the sets of Logging records (e.g. Logging files) for transfer on the IETF CDNI Logging Interface. The objective of this approach is to

reduce the volume of information to be stored and transferred in step [e].

Effect on CDNI interfaces:

- o One additional compression mechanism to be included in the CDNI Logging Interface

Advantages/Drawbacks:

- + No information loss (i.e. all details of each individual chunk delivery are preserved). While this full level of detail may not be needed for some Log consuming applications (e.g. billing), this full level of detail is likely valuable (possibly required) for some Log consuming applications (e.g. debugging)
- + Easier integration (at least in the short term) into existing Logging tools since those are all capable of handling per-chunk records
- + Small extension to CDNI interfaces needed
- + Reduced volume of logging information in step [e]
- + Compression likely to be also applicable to logs for non-HAS content
- High volume of logging information to be handled (storing & processing) at every step of the Logging process from [a] to [g], except [e].

3.4.2.4. Option 4.4: "Full HAS awareness/per-Session-Logs" Approach

In this approach, HAS-awareness is assumed across the CDNs interconnected via CDNI and the necessary information to describe the HAS relationship across all chunks of the same Content Collection is distributed through the CDNI Metadata Interface. In this approach, the dCDN Surrogates leverage the HAS information distributed through the CDNI metadata and their HAS-awareness to generate summarized logging information in the very first place (or alternatively, if per-chunk-logs are generated, to accurately correlate and summarize per-chunk-logs into per-session logs) for exchange over the CDNI Logging interface. The objective of that approach is to operate on lower volume of logging information as early as possible in the successive steps of the Logging process.

Effect on CDNI interfaces:

- o CDNI Metadata Interface: significant extension to convey HAS relationship across chunks of a Content Collection. Note that this extension requires specific support for every HAS-protocol to be supported over the CDNI mesh
- o CDNI Logging Interface: extension to specify summarized per-session logs

Advantages/Drawbacks:

- + Lower volume of logging information to be handled (storing & processing) at every step of the Logging process from [a] to [g]
- + Accurate generation of summarized logs because of HAS awareness in dCDN (for example, where the Surrogate is also serving the Manifest File(s) for a content collection, the Surrogate may be able to extract definitive information about the relationship between all chunks)
- Very significant extensions to CDNI interfaces needed including per HAS-protocol specific support
- Very significant additional requirement for HAS awareness on dCDN and for this HAS-awareness to be consistent with the defined CDNI Logging summarization
- Some information loss (i.e. all details of each individual chunk delivery are not preserved). The actual information loss depends on the summarization approach selected (typically the lower the information loss, the lower the summarization gain) so the right sweet-spot would had ego be selected. While full level of detail may not be needed for some Log consuming applications (e.g. billing), the full level of detail is likely valuable (possibly required) for some Log consuming applications (e.g. debugging)
- Less easy integration (at least in the short term) into existing Logging tools since those are all capable of handling per-chunk records and may not be capable of handling CDNI summarized records
- Challenges in defining behavior (and achieving summarization gain) in the presence of load-balancing of a given HAS-session across multiple Surrogates (in same or different dCDN)

3.4.3. Recommendation

Because of its benefits (in particular simplicity, universal support by CDNs and support by all log-consuming applications), the authors recommend that the per-chunk logging of Option 4.1 be supported by

the CDNI Logging interface as a "High Priority" (as defined in [I-D.[draft-ietf-cdni-requirements](#)]) and be a mandatory capability of CDNs implementing CDNI.

Because of its very low complexity and its benefits in facilitating some useful scenarios (e.g. per-session analytics), we recommend that the CCID mechanisms and Session-ID mechanism of Option 4.2 be supported by the CDNI Metadata interface and the CDNI Logging interface as a "Medium Priority" (as defined in [I-D.[draft-ietf-cdni-requirements](#)]) and be an optional capability of CDNs implementing CDNI.

The authors also recommend that:

- (i) the ability for the uCDN to request that the CCID and Session-ID field be included in log entries provided by the dCDN be supported by the relevant CDNI interfaces, and
- (ii) the ability for the dCDN to include the CCID field and Session-ID in CDNI log entries (when the dCDN is capable of doing so) and indicate so inside the CDNI Logging interface (in line with the "customizable" log format expected to be defined independently of HAS),

be supported as a "Medium Priority" (as defined in [I-D.[draft-ietf-cdni-requirements](#)]) and be an optional capability of CDNs implementing CDNI.

When performing dCDN selection, an uCDN may want to take into account whether a given dCDN is capable of reporting the CCID and session-ID. Thus, the authors recommend that the ability for a dCDN to advertise its support of the optional CCID and SessionID capability be supported by the CDNI request Routing /Footprint and Capabilities Advertisement Interface as a "Medium Priority" (as defined in [I-D.[draft-ietf-cdni-requirements](#)]).

The authors also recommend that a generic mechanism (independent of HAS) be supported allowing the customization of the fields to be reported in logs by CDNs over the CDNI Logging Interface because of the reduction of the logging information volume exchanged across CDNs by removing the information that is not of interest to the other CDN.

Because it can be achieved with very little complexity and it provides some clear storage/communication compression benefits, the authors recommend that, in line with the concept of Option 4.3, some existing very common compression techniques (e.g. gzip) be supported by the CDNI Logging interface as a "Medium Priority" (as defined in [I-D.[draft-ietf-cdni-requirements](#)]) and be an optional capability of

CDNs implementing CDNI.

Because of its complexity, the time it would take to understand the trade-offs of candidate summarization approaches and the time it would take to specify the corresponding support in the CDNI Logging interface, the authors recommend that the log summarization discussed in option 4.4 not be supported by the CDNI Logging interface at this stage and be kept as a candidate topic of great interest for a rechartering of the CDNI WG once the first set of deliverables is produced. When doing so, we suggest to investigate the notion of complementing the "push-style" CDNI logging interface supporting summarization by an on-demand pull-type of interface allowing an uCDN to request the subset of the detailed logging information that it may need but is lost in the summarized pushed information.

The authors note that while a CDN only needs to adhere to the CDNI Logging interface on its external interfaces and can perform logging in a different format within the CDN, any possible CDNI Logging approach effectively places some constraints on the dCDN logging format. For example, to support the "Do-Nothing" Approach, a CDN need to perform and retain per chunk logs. As another example, to support the "Full HAS awareness/per-Session-Logs" Approach, the dCDN cannot operate on logging format that summarize "more than" or "in an incompatible way with" the summarization specified for CDNI Logging. However, the authors feel such constraints are (i) inevitable, (ii) outweighed by the benefits of a standardized logging interface and (iii) acceptable because in case of incompatible summarization, all/most CDNs are capable of reverting to per-chunk logging as per the Do-Nothing Approach that we recommend as the base mandatory approach.

3.5. URL Signing

URL Signing is an authorization method for content delivery. This is based on embedding the HTTP URL with information that can be validated to ensure the request has legitimate access to the content. There are two parts: 1) parameters that convey authorization restrictions (e.g. source IP address and time period) and/or protected URL portion, and 2) message digest that confirms the integrity of the URL and authenticates the URL creator. The authorization parameters can be anything agreed upon between the entity that creates the URL and the entity that validates the URL. A key is used to generate the message digest (i.e. sign the URL) and validate the message digest. The two functions may or may not use the same key.

There are two types of keys used for URL Signing: asymmetric keys and symmetric key. Asymmetric keys always have a key pair made up of a public key and private key. The private key and public key are used

for signing and validating the URL, respectively. A symmetric key is the same key that is used for both functions. Regardless of the type of key, the entity that validates the URL has to obtain the key. Distribution for the symmetric key requires security to prevent others from taking it. Public key can be distributed freely while private key is kept by the URL signer. The method for key distribution is out of scope.

URL Signing operates in the following way. A signed URL is provided by the content owner (i.e. URL signer) to the user during website navigation. When the user selects the URL, the HTTP request is sent to the CDN which validates that URL before delivering the content.

3.5.1. HAS Implications

The authorization lifetime for URL Signing is affected by HAS. The expiration time in the authorization parameters of URL Signing limits the period that the content referenced by the URL can be accessed. This works for URL that directly access the media content. But for HAS content, the Manifest File contains another layer of URL that reference the chunks. The chunk URL that's embedded in the content may be requested at an indeterminate amount of time later. The time period between access to the Manifest File and chunk retrieval may vary significantly. The type of content (i.e. Live or VoD) impacts the time variance as well. HAS content has this property that needs to be addressed for URL Signing.

3.5.2. CDNI Considerations

For CDNI, the two types of request routing are DNS-based and HTTP-based. The use of symmetric vs. asymmetric key for URL Signing has implications on the trust model between CSP and CDNs and the key distribution method that can be used.

DNS-based request routing does not change the URL. In the case of symmetric key, the CSP and the Authoritative CDN have a business relationship that allows them to share a key (or multiple keys) for URL Signing. When the user request a content from the Authoritative CDN, the URL is signed by the CSP. The Authoritative CDN (as a Upstream CDN) redirects the request to a Downstream CDN via DNS. There may be more than one level of redirection to reach the Delivering CDN. The user would obtain the IP address from DNS and send the HTTP request to the Delivering CDN, which needs to validate the URL. This requires the key to be distributed from Authoritative CDN to the Delivering CDN. This may be problematic when the key is exposed to the Delivering CDN that does not have relationship with the CSP. The combination of DNS-based request routing and symmetric key function is a generic issue for URL Signing and not specific to

HAS content. In the case of asymmetric keys, CSP signs URL with its private key. The Delivering CDN validates the URL with the associated public key.

HTTP request routing changes the URL during redirection procedure. In the case of symmetric key, CSP signs the original URL with the same key used by the Authoritative CDN to validate the URL. The Authoritative CDN (as a Upstream CDN) redirects the request to the Downstream CDN. The new URL is signed by the Upstream CDN with the same key used by the Downstream CDN to validate that URL. The key used by the Upstream CDN to validate the original URL is expect to be different than the key used to sign the new URL. In the case of asymmetric keys, CSP signs the original URL with its private key. Authoritative CDN validates that URL with the CSP's public key. The Authoritative CDN redirects the request to the Downstream CDN. The new URL is signed by the Upstream CDN with its private key. The Downstream CDN validates that URL with the Upstream CDN's public key. There may be more than one level of redirection to reach the Delivering CDN. The URL Signing operation described previously applies at each level between the Upstream CDN and Downstream CDN for both the symmetric key and asymmetric keys cases.

URL Signing requires support in most of the CDNI Interfaces. The CDNI Metadata interface should specify the content that is subject to URL signing and provide information to perform the function. The Downstream CDN should inform the Upstream CDN that it supports URL Signing in the asynchronous capabilities information advertisement as part of the Request Routing interface. This allows the CDN selection function in request routing to choose the Downstream CDN with URL signing capability when the CDNI metadata of the content requires this authorization method. The Logging interface provides information on the authorization method (e.g. URL Signing) and related authorization parameters used for content delivery. Having the information in the URL is not sufficient to know that the surrogate enforced the authorization. URL Signing has no impact on the Control interface.

3.5.3. Option 5.1: Do Nothing

"Do Nothing" approach means that CSP can only perform URL Signing for the top level Manifest File. The top level Manifest File contains chunk URLs or lower level Manifest File URLs, which are not modified (i.e. no URL Signing for the embedded URLs). In essence, the lower level Manifest Files and chunks are delivered without content access authorization.

Effect on CDNI interfaces:

- o None

Advantages/Drawbacks:

- + Top level Manifest File access is protected
- + Upstream CDN and Downstream CDN do not need to be aware of HAS content
- Lower level Manifest Files and chunks are not protected, making this approach unqualified for content access authorization

3.5.4. Option 5.2: Flexible URL Signing by CSP

In addition to URL Signing for the top level Manifest File, CSP performs flexible URL Signing for the lower level Manifest Files and chunks. For each HAS session, the top level Manifest File contains signed chunk URLs or signed lower level Manifest File URLs for the specific session. The lower level Manifest File contains session-based signed chunk URLs. CSP generates the Manifest Files dynamically for the session. The chunk (segment/fragment) is delivered with content access authorization using flexible URL Signing which protects the invariant portion of the URL. Segment URL (e.g. HLS) is individually signed for the invariant URL portion (Relative URL) or the entire URL (Absolute URL without Redirection) in the Manifest File. Fragment URL (e.g. Smooth Streaming) is signed for the invariant portion of the template URL in the Manifest File. More details are provided later in this section. The URL Signing expiration time for the chunk needs to be long enough to play the video. There are implications of signing the URLs in the Manifest File. For Live content, the Manifest Files are requested at a high frequency. For VoD content, the Manifest File may be quite large. URL Signing can add more computational load and delivery latency in high volume cases.

For HAS content, the Manifest File contains the Relative Locator, Absolute Locator without Redirection, or Absolute Locator with Redirection for specifying the chunk location. Signing the chunk URL requires CSP to know the portion of the URL that remains when the content is requested from the Delivery CDN surrogate.

For Absolute URL without Redirection, the CSP knows that the chunk URL which is explicitly linked with the delivery CDN surrogate and can sign the URL based on that information. Since the entire URL is set and does not change, the surrogate can validate the URL. The CSP and the Delivery CDN are expected to have a business relationship in this case. So either symmetric key or asymmetric keys can be used for URL Signing.

For Relative URL, the URL of the Manifest File provides the root location. The method of request routing affects the URL used to ultimately request the chunk from the Delivery CDN surrogate. For DNS, the original URL does not change. This allows CSP to sign the chunk URL based on the Manifest File URL and the Relative URL. For HTTP, the URL changes during redirection. In this case, CSP does not know the redirected URL that will be used to request the Manifest File. This uncertainty makes it impossible to accurately sign the chunk URLs in the Manifest File. Basically, URL Signing using this reference method, "as is" for entire URL protection, is not supported. However, instead of signing the entire URL, the CSP signs the Relative URL (i.e. invariant portion of the URL) and conveys the protected portion in the authorization parameters embedded in the chunk URL. This approach works the same way as Absolute URL without Redirection, except the HOST part and (part of) the PATH part of the URL are not signed and validated. The security level should remain the same as content access authorization ensures that the user that requested the content has the credentials. This scheme does not seem to compromise the authorization model since the resource is still protected by the authorization parameters and message digest. Perhaps, further evaluation on security would be helpful.

For Absolute URL with Redirection, the method of request routing affects the URL used to ultimately request the chunk from the Delivery CDN surrogate. This case has the same conditions as the Relative URL. The difference is that the URL is for the chunk instead of the Manifest File. For DNS, the chunk URL does not change and can be signed by the CSP. For HTTP, the URL used to deliver the chunk is unknown to the CSP. In this case, CSP cannot sign the URL and this method of reference for the chunk is not supported.

Effect on CDNI interfaces:

- o Requires the ability to exclude the variant portion of URL in the signing process (NOTE: Issue is specific to URL Signing support for HAS content and not CDNI?)

Advantages/Drawbacks:

- + Manifest File and chunks are protected
- + Upstream CDN and Downstream CDN do not need to be aware of HAS content
- + DNS-based request routing with asymmetric keys and HTTP-based request routing for Relative URL and Absolute URL without Redirection works

- CSP has to generate Manifest Files with session-based signed URLs and becomes involved in content access authorization for every HAS session
- Manifest Files are not cacheable
- DNS-based request routing with symmetric key may be problematic due to need for transitive trust between CSP and Delivery CDN
- HTTP-based request routing for Absolute URL with Redirection does not work because the URL used Delivery CDN surrogate is unknown to the CSP

3.5.5. Option 5.3: Flexible URL Signing by Upstream CDN

This is similar to the previous section, with the exception that the Upstream CDN performs flexible URL for the lower level Manifest Files and chunks. URL Signing for the top level Manifest File is still provided by the CSP.

Effect on CDNI interfaces:

- o Requires the ability to exclude the variant portion of URL in the signing process (NOTE: Issue is specific to URL Signing support for HAS content and not CDNI?)

Advantages/Drawbacks:

- + Manifest File and chunks are protected
- + CSP does not need to be involved in content access authorization for every HAS session
- + Downstream CDN does not need to be aware of HAS content
- + DNS-based request routing with asymmetric keys and HTTP-based request routing for Relative URL and Absolute URL without Redirection works
- Upstream CDN has to generate Manifest Files with session-based signed URLs and becomes involved in content access authorization for every HAS session
- Manifest Files are not cacheable

- Manifest File needs to be distributed through the uCDN
- DNS-based request routing with symmetric key may be problematic due to need for transitive trust between uCDN and non-adjacent Delivery CDN
- HTTP-based request routing for Absolute URL with Redirection does not work because the URL used Delivery CDN surrogate is unknown to the uCDN

3.5.6. Option 5.4: Authorization Group ID and HTTP Cookie

Based on the Authorization Group ID metadata, CDN validates the URL Signing or validates the HTTP cookie for request of content in the group. CSP performs URL Signing for the top level Manifest File. The top level Manifest File contains lower level Manifest File URLs or chunk URLs. The lower level Manifest Files and chunks are delivered with content access authorization using HTTP cookie that contains session state associated with authorization of the top level Manifest File. The Group ID Metadata is used to associate the related content (i.e. Manifest Files and chunks). It also specifies content (e.g. regexp method) that needs to be validated by either URL Signing or HTTP cookie. Note that the creator of the metadata is HAS-aware. Duration of the chunk access may be included in the URL Signing of the top level Manifest File and set in the cookie. Alternatively, the access control duration could be provided by the CDNI Metadata interface.

Effect on CDNI interfaces:

- o CDNI Metadata Interface - Authorization Group ID metadata identifies the content that is subject to validation of URL Signing or validation of HTTP cookie associated with the URL Signing
- o CDNI Logging Interface - Report the authorization method used to validate the request for content delivery

Advantages/Drawbacks:

- + Manifest File and chunks are protected
- + CDN does not need to be aware of HAS content
- + CSP does not need to change the Manifest Files

- Authorization Group ID metadata is required (i.e. CDNI Metadata Interface enhancement)
- Requires the use of HTTP cookie which may not be acceptable in some environments (e.g. where some targeted User-Agents do not support HTTP Cookie)
- Manifest File has to be delivered by surrogate

3.5.7. Option 5.5: HAS-awareness with HTTP Cookie in CDN

CDN is aware of HAS content and uses URL Signing and HTTP cookie for content access authorization. URL Signing is fundamentally about authorizing access to a Content Item or its specific Content Collections (representations) for a specific user during a time period with possibly some other criteria. A chunk is an instance of the sets of chunks referenced by the Manifest File for the Content Item or its specific Content Collections. This relationship means that once the Downstream CDN has authorized the Manifest File, it can assume that the associated chunks are implicitly authorized. The new function for the CDN is to link the Manifest File with the chunks for the HTTP session. This can be accomplished by using an HTTP cookie for the HAS session.

After validating the URL and detecting that the requested content is a top level Manifest File, the delivery CDN surrogate sets a HTTP cookie with a signed session token for the HTTP session. When a request for a lower level Manifest File or chunk arrives, the surrogate confirms that the HTTP cookie value contains the correct session token. If so, the lower level Manifest File or chunk is delivered due to transitive authorization property. Duration of the chunk access may be included in the URL Signing of the top level Manifest File and set in the cookie. The details of the operation are left to be determined later.

Effect on CDNI interfaces:

- o CDNI Metadata Interface - New metadata identifies the content that is subject to validation of URL Signing and information in the cookie for the type of HAS content
- o Request Routing interface - Downstream CDN should inform the Upstream CDN that it supports URL Signing for known HAS content types in the asynchronous capabilities information advertisement. This allows the CDN selection function in request routing to choose the appropriate Downstream CDN when the CDNI metadata identifies the content

- o CDNI Logging Interface - Report the authorization method used to validate the request for content delivery

Advantages/Drawbacks:

- + Manifest File and chunks are protected
- + CSP does not need to change the Manifest Files
- Requires full HAS awareness on part of Upstream CDN and Downstream CDN
- Requires CDNI Interfaces extensions
- Requires the use of HTTP cookie which may not be acceptable in some environments (e.g. where some targeted User-Agents do not support HTTP Cookie)
- Manifest File has to be delivered by surrogate

3.5.8. Option 5.6: HAS-awareness with Manifest in CDN

CDN is aware of HAS content and uses URL Signing for content access authorization of Manifest File and chunk. CDN generates or rewrites the Manifest Files and learns about the chunks based on the Manifest File. The embedded URLs in the Manifest File are signed by the CDN. Duration of the chunk access may be included in the URL Signing. The details of the operation are left to be determined later. Since this approach is based on signing the URLs in the Manifest File, the implications for Live and VoD content mentioned in [Section 3.5.4](#) apply.

Effect on CDNI interfaces:

- o CDNI Metadata Interface - New metadata identifies the content that is subject to validation of URL Signing and information in the cookie for the type of HAS content
- o Request Routing interface - Downstream CDN should inform the Upstream CDN that it supports URL Signing for known HAS content types in the asynchronous capabilities information advertisement. This allows the CDN selection function in request routing to choose the appropriate Downstream CDN when the CDNI metadata identifies the content
- o CDNI Logging Interface - Report the authorization method used to validate the request for content delivery

Advantages/Drawbacks:

- + Manifest File and chunks are protected
- + CSP does not need to change the Manifest Files
- Requires full HAS awareness on part of Upstream CDN and Downstream CDN
- Requires CDNI Interfaces extensions
- Requires CDN to generate or rewrite the Manifest File
- Manifest File has to be delivered by surrogate

3.5.9. Recommendation

The authors consider that Option 5.1 (Do Nothing) is not suitable for access control of HAS content.

Where the HTTP Cookie mechanism is supported by the targeted User-Agents and the security requirements can be addressed through proper use of HTTP Cookies, the authors recommend use of Option 5.4 (Authorization Group ID with HTTP Cookie) and therefore be supported by the CDNI solution. This method does not require manifest file manipulation which may be a significant deployment obstacle. Otherwise, the authors recommend that Option 5.2 (Flexible URL Signing by the CSP) or Option 5.3 (Flexible URI Signing by the Upstream CDN) be used and therefore that flexible URI be supported by the CDNI solution. Option 5.2 and Option 5.3 protect all the content, does not require Downstream CDN to be aware of HAS, does not impact CDNI interfaces, supports all different types of devices, and supports the common cases of request routing for HAS content (i.e. DNS-based request routing with asymmetric keys and HTTP-based request routing for Relative URL).

HAS-awareness in CDN (Option 5.5 and Option 5.6) have some advantages that should be considered for future support (e.g. CDN that is aware of HAS content can manage the content more efficiently at a broader context. Content distribution, storage, delivery, deletion, access authorization, etc. can all benefit.). Including HAS-awareness as part of the current CDNI charter, however, would almost certainly delay the CDNI WG's milestones, and the authors therefore do not recommend it right now.

3.6. Content Purge

At some point in time, a uCDN might want to remove content from a dCDN. With regular content, this process can be relatively straightforward; a uCDN will typically send the request for content removal to the dCDN including a reference to the content which it wants to remove (e.g. in the form of a URL). Due to the fact that HAS content consists of large groups of files however, things might be more complex. [Section 3.1](#) describes a number of different scenarios for doing file management on these groups of files, while [Section 3.2](#) list the options for performing Content Acquisition on these Content Collections. This section will present the options for requesting a Content Purge for the removal of a Content Collection from a dCDN.

3.6.1. Option 6.1: No HAS awareness

The most straightforward way to signal content purge requests is to just send a single purge request for every file that makes up the Content Collection. While this method is very simple and does not require HAS awareness, it obviously creates a signaling overhead between the uCDN and dCDN since a reference is to be provided for each content chunk to be purged.

Effect on CDNI interfaces:

- o None

Advantages/Drawbacks (apart from those listed under Option 3.3):

- + Does not require changes to the CDNI Interfaces or HAS awareness
- Requires individual purge request for every file making up a Content Collection (or, alternatively, requires the ability to convey references to all the chunks making up a Content Collection inside a purge request) which creates signaling overhead

3.6.2. Option 6.2: Purge Identifiers

There exists a potentially more efficient method for performing content removal of large numbers of files simultaneously. By including a "Purge Identifier (Purge-ID)" in the metadata of a particular file, it is possible to virtually group together different files making up a Content Collection. A Purge-ID can take the form of an arbitrary number or string which is communicated as part of the CDNI Metadata Interface and which is the same for all files making up a particular Content Item, and different across different Content Items. If a uCDN wants to request the dCDN to remove a Content

Collection, it can send a purge request containing this Purge-ID. The dCDN can then remove all files that share the corresponding Purge-ID.

The advantage of this method is that it is relatively simple to use by both the dCDN and uCDN and requiring only limited additions to the CDNI Metadata Interface and CDNI Control Interface.

The Purge-ID is similar to the Content Collection ID discussed in [Section 3.4.2.2](#) for handling HAS Logging and we note that further thought is needed to determine whether the CCID and Purge-ID should be collapsed into a single element or remain separate elements.

Effect on CDNI interfaces:

- o CDNI Metadata Interface: Add metadata field for indicating Purge-ID
- o CDNI Control Interface: Add functionality to convey a Purge-ID in purge requests

Advantages/Drawbacks:

- + Allows for efficient purging of content from a dCDN
- + Does not require HAS awareness on part of dCDN

[3.6.3.](#) Recommendation

Based on the listed pros and cons, the authors recommend the WG to have mandatory support Option 1.1, the 'Do Nothing'-approach. In addition, because of its very low complexity and its benefit in facilitating low-overhead purge of large numbers of content items simultaneously, the authors recommend that the Purge IDentifier of Option 6.2 be supported as an optional feature by the CDNI Metadata interface and the CDNI Control interface.

[3.7.](#) Other issues

This section includes some HAS-specific issues that came up during the discussion of this draft and which do not fall under any of the categories discussed in the previous sections.

- As described in [Section 2.2](#), a Manifest File might either be delivered by a CDN or by the CSP, thereby being invisible to the CDN delivering the chunks. Obviously, the decision on whether the CDN or CSP delivers the Manifest File is made between the uCDN and CSP, and the dCDN has no choice in the matter. However, some

dCDNs might only want to offer their services in the cases where they have access to the Manifest File (e.g. because their internal architecture is based around the knowledge inside the Manifest File). For these cases, it might be useful to include a field in the CDNI Capability Advertisement to allow dCDNs to advertise the fact that they require access to the Manifest File.

4. IANA Considerations

This document makes no request of IANA.

5. Security Considerations

This document does not discuss security issues around HTTP or HAS delivery. Those are expected to be discussed in the CDNI WG documents including [[I-D.ietf-cdni-framework](#)].

6. Acknowledgements

The authors would like to thank Kevin Ma, Stef van der Ziel, Bhaskar Bhupalam, Mahesh Viveganandhan, Larry Peterson, Ben Niven-Jenkins and Matt Caulfield for their valuable contributions to this document.

7. References

[7.1. Normative References](#)

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

[7.2. Informative References](#)

[I-D.[draft-bertrand-cdni-logging](#)]
Bertrand, G., Ed. and E. Stephan, "CDNI Logging Interface".

[I-D.[draft-ietf-cdni-requirements](#)]
Leung, K. and Y. Lee, "Content Distribution Network Interconnection (CDNI) Requirements, [draft-ietf-cdni-requirements-03](#)", June 2012.

[I-D.[ietf-cdni-framework](#)]
Peterson, L. and B. Davie, "Framework for CDN

Interconnection", [draft-ietf-cdni-framework-02](#) (work in progress), December 2012.

[RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax, [RFC3986](#)", January 2005.

Authors' Addresses

Ray van Brandenburg
TNO
Brassersplein 2
Delft 2612CT
the Netherlands

Phone: +31-88-866-7000
Email: ray.vanbrandenburg@tno.nl

Oskar van Deventer
TNO
Brassersplein 2
Delft 2612CT
the Netherlands

Phone: +31-88-866-7000
Email: oskar.vandeventer@tno.nl

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

Kent Leung
Cisco Systems
170 West Tasman Drive
San Jose, CA 95134
USA

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