

Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 14, 2013

F. Le Faucheur
M. Viveganandhan
K. Leung
Cisco
July 13, 2012

CDNI Logging Formats for HTTP and HTTP Adaptive Streaming Deliveries
draft-lefaucheur-cdni-logging-delivery-01

Abstract

The interconnection of Content Distribution Networks (CDNs) is motivated by several use cases. CDN Interconnection can be achieved through four CDNI interfaces, one of which is the CDNI Logging interface. This document discusses CDNI logging formats for content deliveries performed using HTTP or HTTP adaptive streaming.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 14, 2013.

Copyright Notice

Copyright (c) 2012 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	3
1.1.	Terminology	4
2.	CDNI Logging for regular HTTP Delivery	4
2.1.	Regular HTTP Log Triggers	4
2.2.	Regular HTTP Log Fields	4
3.	CDNI Logging for HTTP Adaptive Streaming	6
3.1.	HAS Chunk-Based Log Triggers	6
3.2.	HAS Chunk-Based Log fields	6
4.	Performance Monitoring	7
5.	CDNI Log Encoding and Transport	7
5.1.	CDNI Log Header Information	8
5.2.	CDNI per-Log-Record Information	8
5.3.	CDNI Log Footer Information	9
5.4.	CDNI Customized Log Format	9
6.	Impact on CDNI Metadata Requirements	9
7.	Impact on CDNI Footprint and Capabilities Advertisement . . .	10
8.	Generation of CDNI Logs	10
9.	IANA Considerations	10
10.	Security Considerations	11
11.	Acknowledgements	11
12.	References	11
12.1.	Normative References	11
12.2.	Informative References	11
	Authors' Addresses	12

1. Introduction

The interconnection of Content Distribution Networks (CDNs) is motivated by several use cases, such as those described in [\[I-D.ietf-cdni-use-cases\]](#). The overall problem space for CDN Interconnection (CDNI) is described in [\[I-D.ietf-cdni-problem-statement\]](#). The CDN Interconnection framework is defined in [\[I-D.ietf-cdni-framework\]](#) and the requirements for the CDN Interconnection solution are specified in [\[I-D.ietf-cdni-requirements\]](#).

One of the CDNI interfaces defined in these documents is the CDNI Logging interface whose description is quoted here:

"o CDNI Logging interface: This interface allows the Logging system in interconnected CDNs to communicate the relevant activity logs in order to allow log consuming applications to operate in a multi-CDN environments. For example, an upstream CDN may collect delivery logs from a downstream CDN in order to perform consolidated charging of the CSP or for settlement purposes across CDNs. Similarly, an upstream CDN may collect delivery logs from a downstream CDN in order to provide consolidated reporting and monitoring to the CSP."

The logging interface is discussed in details in [\[I-D.bertrand-cdni-logging\]](#). The present document identifies a number of additional considerations regarding logging for content delivery performed using HTTP.

[\[I-D.brandenburg-cdni-has\]](#) discusses the interactions of HTTP Adaptive Streaming with CDNI Logging and provides recommendations on how to specifically perform CDNI Logging for delivery performed using HTTP Adaptive Streaming. The present document discusses in more details how these recommendations would impact the CDNI Logging interface.

[Section 2](#) discusses the aspect of CDNI Logging that are specific to content delivery performed via regular HTTP. [Section 3](#) discusses the aspect of CDNI Logging that are specific to content delivery performed via HTTP adaptive streaming. [Section 5](#) discusses the aspect of CDNI Logging that are generic to all the CDNI Logs for delivery.

CDNI Logging for other events than content delivery (e.g. failures, request routing, service monitoring,...) are not discussed in the present document.

1.1. Terminology

This document uses all terminology defined in [\[I-D.brandenburg-cdni-has\]](#).

2. CDNI Logging for regular HTTP Delivery

This section discusses the triggers on which a CDNI log is generated for delivery of content using regular HTTP delivery, as well as the fields contained in the corresponding log.

2.1. Regular HTTP Log Triggers

A CDNI Log is generated for a regular HTTP delivery on the following triggers:

Event	Description
content Request	Reception and processing of a request for a content

2.2. Regular HTTP Log Fields

A CDNI Log for regular HTTP delivery contains the following fields (or a subset thereof as discussed in [Section 5.4](#)):

Field	Description	Examples
Current-Time	Time, in milliseconds, at which the request was received.	[20/Feb/2012:00:29.510+0200]
Time-to-Serve	Time, in microseconds, taken to complete the request.	952195

Client-IP	IP address of the requesting client.	203.0.113.2
Action	Squid action describing how the request was treated locally (e.g., cache hit/miss)	TCP_HIT, TCP_MISS, ...
Status-Returned	HTTP response code.	200, 404, ...
Bytes-Transferred	Bytes sent to the client, including the headers.	23567992
Method	HTTP request method.	GET
URI	URI of requested content.	http://cache3.cdn1.com/movie/ice/icemovie.mpg
Content-Type	MIME-Type from the reply header.	video/mpeg
User-Agent (and possibly some other HTTP headers)	content of the User-Agent HTTP Header	Mozilla/5.0 (X11; Linux x86_64; rv:12.0) Gecko/20100101 Firefox/12.0

URI-Signing-Vali	Flag	0/1	
dation	indicating		
	whether		
	URI		
	signature		
	validation		
	was		
	performed		
+-----+	+-----+	+-----+	+-----+

3. CDNI Logging for HTTP Adaptive Streaming

This section discusses the CDNI logs for delivery of content using HTTP adaptive streaming. In line with the recommendations of [\[I-D.brandenburg-cdni-has\]](#), this document proposes that a log record be generated for delivery of each chunk or manifest file.

3.1. HAS Chunk-Based Log Triggers

A chunk-based Log record is generated for HAS delivery on the following triggers:

+-----+	+-----+	+-----+
Event	Description	
+-----+	+-----+	+-----+
Chunk/Manifest	Reception of a request for a Segment (or	
Request	Manifest File)	
+-----+	+-----+	+-----+

3.2. HAS Chunk-Based Log fields

A chunk-based Log record for HTTP adaptive streaming may contain all the same fields as a CDNI Log for a regular HTTP request, as well as the following additional fields:

+-----+	+-----+	+-----+
Field	Description	Examples
+-----+	+-----+	+-----+
Content	Identifier for	format and scope of unicity are
Collection	Content	TBD
ID	Collection	

Session-id	a string	6141F5795BE774691D234A0465B9667A	
	generated by the		
	delivering CDN		
	and unique (to		
	the delivering		
	CDN) to identify		
	the Session. (*)		
+-----+-----+-----+-----+			

(*) 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 mechanism to synchronize session-IDs across Surrogates.

4. Performance Monitoring

The CDNI Log fields listed in the previous sections allow monitoring of essential delivery performance indicators across the CDN Mesh. For example, for a regular HTTP delivery, these fields allow tracking of the time taken to serve the requested content, tracking of content delivery failures, tracking of partial deliveries and tracking of cache hit ratios. As another example, for HTTP adaptive streaming, these fields allow tracking of the presentation (and its fluctuation over time) served to the End-User as well as the End-User random content access (e.g. Play/Stop/Seek).

Subsequent versions of this document may discuss potential additional log fields for enhanced performance monitoring.

5. CDNI Log Encoding and Transport

Details for CDNI Log encoding and transport will be specified in subsequent versions. WE observe that this is expected to allow optional use of common compression techniques (e.g. gzip). However,

[Section 5.1](#) identifies the information that is to be included in the header of CDNI Logs, [Section 5.2](#) identifies information that is to be attached to every CDNI Log record and [Section 5.3](#) identifies the information that is to be included in the footer of CDNI Logs. Finally, [Section 5.4](#) discusses the notion of customized Logging.

5.1. CDNI Log Header Information

The header of CDNI Logs contains the following fields:

Field	Description	Examples
Format-Version	Version of the CDNI Log format.	v1.0
Log-Field-List	The list of the fields provided in the log records	time cs-method cs-uri
Log-ID	Unique identifier for the CDNI Log (facilitates detection of duplicate Logs and tracking in case of aggregation).	
Log-Timestamp	Time, in milliseconds, the CDNI Log was generated.	[20/Feb/2012:00:29.510+0200]

5.2. CDNI per-Log-Record Information

In addition to the log fields discussed in previous sections, each CDNI log record contains the following fields:

Field	Description	Examples
Log Record Digest	Digest of the Log records (facilitates recovery of uncorrupted Log records inside a corrupted CDNI Log)	

5.3. CDNI Log Footer Information

The footer of CDNI Logs contains the following fields:

Field	Description	Examples
Log Digest	Digest of the complete Log (facilitates detection of Log corruption)	

5.4. CDNI Customized Log Format

This document proposes that customized logs be supported by CDNI in the following manner:

- o the uCDN uses the CDNI Metadata interface to indicate to the dCDN which subset of the CDNI logging fields are to be provided in a log record for corresponding to a request for a given content
- o the dCDN provides, via the CDNI Logging interface, log records containing the subset of CDNI logging fields requested by the uCDN.
- o The dCDN explicitly lists in the CDNI Log Header the fields actually provided (as discussed in [Section 5.1](#)).

6. Impact on CDNI Metadata Requirements

We request that the following requirements be added in section 6 of [\[I-D.ietf-cdni-requirements\]](#):

"

META-X [HIGH] The CDNI Metadata Distribution interface shall support an OPTIONAL mechanism allowing the Upstream CDN to indicate to the Downstream CDN which CDNI Log fields are to be provided for all, for specific sets of, or for specific content items delivered using HTTP. A CDNI implementation that does not support this optional CDNI Metadata Distribution Interface mechanism MUST ignore this log format indication and generate CDNI logging format for adaptive streaming using the default set of CDNI Logging fields. [Editor's note: A default set of logging fields need to be defined]

META-X [MID] The CDNI Metadata Distribution interface shall allow the uCDN to signal to the dCDN the Content Collection ID value for all, for specific sets of, or for specific content items delivered using

HTTP. Whenever the dCDN is instructed by the uCDN (using the customized logging mechanism described in [Section 5.4](#)) to report the Content Collection ID field in the log records, the dCDN is to use the value provided through the CDNI Metadata interface for the corresponding content.

"

7. Impact on CDNI Footprint and Capabilities Advertisement

We request that the following requirement be added in section 5 of [\[I-D.ietf-cdni-requirements\]](#):

"

REQ-X [MID] The CDNI Request Routing/Footprint and Advertisement Interface shall support advertisement of the following capabilities:

- o support for customized CDNI Logging
- o support of Content Collection ID logging
- o support for Session-ID logging

"

8. Generation of CDNI Logs

Like other CDNI interfaces, the CDNI Logging interface specifies operations across CDNs and not inside a CDN. Therefore, the log formats specified in this document apply to CDNI logging information exchanged across CDNs and does not constrain the process for generating such inter-CDN logs within a given CDN. The format of the logs generated by a given CDN Surrogate is beyond the scope of the present document. We observe that a given CDN could elect to have its Surrogates natively generate logs in the same format as the one to be used for exchange with another CDN, or that the CDN could elect to have its Surrogates generate logs in any other format (as long as they include the necessary information) and have these logs then reformatted prior to exchange with another CDN.

9. IANA Considerations

[This will be specified in subsequent versions].

10. Security Considerations

CDNI Logs exchanged over the CDNI Logging interface can be consumed by very sensitive applications including inter-CDN accounting and billing. The associated security concerns are discussed in [[I-D.ietf-cdni-framework](#)]. At this stage, we observe that the CDNI Logging interface can leverage the existing security mechanisms supported by the underlying transport protocol that will be selected for transport of CDNI Logs (e.g. to support authentication of the entities exchanging CDNI Logs through the CDNI Logging interface, to support privacy and protection against spoofing). This will be further discussed in subsequent versions of this document.

11. Acknowledgements

The authors want to thank Gene Halbrooks for his input into this document.

12. References

12.1. Normative References

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

12.2. Informative References

- [I-D.bertrand-cdni-logging]
Bertrand, G. and S. Emile, "CDNI Logging Interface", [draft-bertrand-cdni-logging-00](#) (work in progress), February 2012.
- [I-D.brandenburg-cdni-has]
Brandenburg, R., Deventer, O., Faucheur, F., and K. Leung, "Models for adaptive-streaming-aware CDN Interconnection", [draft-brandenburg-cdni-has-03](#) (work in progress), July 2012.
- [I-D.ietf-cdni-framework]
Peterson, L. and B. Davie, "Framework for CDN Interconnection", [draft-ietf-cdni-framework-00](#) (work in progress), April 2012.
- [I-D.ietf-cdni-problem-statement]
Niven-Jenkins, B., Faucheur, F., and N. Bitar, "Content Distribution Network Interconnection (CDNI) Problem

Statement", [draft-ietf-cdni-problem-statement-08](#) (work in progress), June 2012.

[I-D.ietf-cdni-requirements]

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

[I-D.ietf-cdni-use-cases]

Bertrand, G., Emile, S., Burbridge, T., Eardley, P., Ma, K., and G. Watson, "Use Cases for Content Delivery Network Interconnection", [draft-ietf-cdni-use-cases-09](#) (work in progress), July 2012.

Authors' Addresses

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

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

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

Email: mvittal@cisco.com

Kent Leung
Cisco Systems
3625 Cisco Way
San Jose 95134
USA

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

