

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 24, 2014

M. Caulfield
Cisco Systems
October 21, 2013

CDNI Rate Pacing
draft-caulfield-cdni-rate-pacing-00

Abstract

Rate pacing is a class of network traffic shaping which limits the transmission rate of data over a network. This document defines CDNI extensions for downstream CDNs to support rate pacing on behalf of upstream CDNs.

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 April 24, 2014.

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	2
2. Rate Pacing Algorithm	3
3. CDNI Interfaces Impact	3
3.1. Footprint & Capabilities Interface	3
3.2. Request Routing Redirection Interface	3
3.3. Metadata Interface	4
3.3.1. RatePacing Metadata	4
3.4. Logging Interface	4
3.5. Control Interface	5
4. IANA Considerations	5
5. Security Considerations	5
6. Acknowledgements	5
7. Normative References	5
Author's Address	6

1. Introduction

Rate pacing is a class of network traffic shaping which limits the transmission rate of data over a network. In the context of a Content Delivery Network (CDN), rate pacing provides an important business advantage to a Content Service Provider (CSP) by ensuring that a CDN which is delivering content on behalf of that CSP does not deliver significantly more data than necessary to an end client.

For example, suppose an end client is watching some Constant Bit Rate (CBR) video encoded at 1500 kbps. In the absence of rate pacing, the CDN delivering this content may send it to the client at 3000 kbps. If the client chooses to terminate the session before watching the entire video, up to half the transmitted data is wasted. This waste leads to unnecessary cost for the CSP and diminished useful capacity for the CDN.

Rate pacing requires configuration on a per-content basis. In order to enable rate pacing in a CDNI environment, the CDNI interfaces need to be extended to optionally support this feature.

This document describes:

1. a rate pacing algorithm for CDNs

2. CDNI interface extensions required for supporting rate pacing

2. Rate Pacing Algorithm

A token bucket algorithm SHOULD be used for implementing rate pacing. Other algorithms MAY be used but the traffic shape MUST fit within the same envelope of a token bucket algorithm. Token bucket is described by [RFC1363].

The token bucket algorithm is characterized by two parameters:

1. Rate - the number of tokens added to the bucket per second
2. Size - the maximum number of tokens in the bucket

For the purpose of this document, each token represents one byte.

3. CDNI Interfaces Impact

3.1. Footprint & Capabilities Interface

[I-D.ietf-cdni-footprint-capabilities-semantics] defines the CDNI Footprint and Capabilities semantics. But at the time of writing, no FCI syntax specification has been accepted as a working group document.

[I-D.ietf-cdni-footprint-capabilities-semantics] states that:

"The CDNI FCI specification SHOULD define the registry (and the rules for adding new entries to the registry) for the different capability types. Each capability type MAY further have a list of valid values. The individual CDNI interface specifications which define a given capability SHOULD define any necessary registries (and the rules for adding new entries to the registry) for the values advertised for a given capability type."

This document defines a new capability type: "RatePacing" to be added to the FCI capability types registry. No value needs to be advertised for this capability and therefore no additional value registry need to be defined.

A CDN MAY advertise the "RatePacing" capability in the FCI if it implements this specification.

3.2. Request Routing Redirection Interface

A new RatePacing metadata object is defined to represent the configuration for rate pacing. The RatePacing object has MIME type

"application/cdni.RatePacing". RatingPacing MAY appear within the metadata list of either HostMetadata or PathMetadata (i.e. may have either host-level scope or a path-level scope). The following section defines the properties of the RatingPacing object.

3.3. Metadata Interface

A new RatePacing metadata object is defined to represent the configuration for rate pacing. The RatePacing object has MIME type "application/cdni.RatePacing". RatingPacing MAY appear within the metadata list of either HostMetadata or PathMetadata (i.e. may have either host-level scope or a path-level scope). The following section defines the properties of the RatingPacing object.

3.3.1. RatePacing Metadata

The presence of the RatePacing Metadata indicates that a dCDN MUST comply with the Rate Pacing Algorithm defined by this specification in order to deliver a piece of content.

Property: rate

Description: Rate of tokens per second to be added to the bucket as described by the token bucket algorithm. This value MUST be a positive integer. Each token represents one byte.

Type: Integer

Mandatory-to-Specify: Yes.

Property: size

Description: Maximum number of tokens per bucket as described by the token bucket algorithm. This value MUST be a positive integer.

Type: Integer

Mandatory-to-Specify: Yes.

3.4. Logging Interface

The rate at which a piece of content was delivered MAY be indicated via the LI. The "sc-rate" field indicates the rate in bytes per second as a decimal number. The bytes measured should correspond to the sc-entity-bytes field.

sc-rate:

format: DEC

field value: the average rate in bytes per second at which a response was delivered from Surrogate to client.

occurrence: there MUST be zero or exactly one instances of this field.

3.5. Control Interface

The CI is not impacted by rate pacing.

4. IANA Considerations

This document requests the following of IANA:

Addition of RatePacing in the CDNI Capability Registry defined in ???.

Addition of RatePacing in the CDNI GenericMetadata Type Registry defined in [I-D.ietf-cdni-metadata].

Addition of sc-rate in the CDNI Logging Field Names Registry defined in [I-D.ietf-cdni-logging].

5. Security Considerations

A malicious CSP might attempt to use rate pacing to instruct a dCDN to delivery some content at a very high rate thereby amplifying a DDOS attack. The decision to enforce a rate is left to the discretion of a dCDN. An implementation of rate pacing should implement reasonable upper bound to avoid such cases.

6. Acknowledgements

The author would like to thank Francois Le Faucheur for his contributions and feedback.

7. Normative References

[I-D.ietf-cdni-footprint-capabilities-semantics]
Seedorf, J., Peterson, J., Previdi, S., Brandenburg, R.,
and K. Ma, "CDNI Request Routing: Footprint and
Capabilities Semantics", draft-ietf-cdni-footprint-
capabilities-semantics-00 (work in progress), July 2013.

[I-D.ietf-cdni-logging]

Faucheur, F., Bertrand, G., Oprescu, I., and R.
Peterkofsky, "CDNI Logging Interface", draft-ietf-cdni-
logging-08 (work in progress), October 2013.

[I-D.ietf-cdni-metadata]

Niven-Jenkins, B., Murray, R., Watson, G., Caulfield, M.,
Leung, K., and K. Ma, "CDN Interconnect Metadata", draft-
ietf-cdni-metadata-02 (work in progress), July 2013.

[RFC1363] Partridge, C., "A Proposed Flow Specification", RFC 1363,
September 1992.

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

Author's Address

Matt Caulfield
Cisco Systems
1414 Massachusetts Ave
Boxborough, MA 01719
USA

Phone: +1 978 936 9307
Email: mcaulfie@cisco.com

Network Working Group
Internet-Draft
Intended Status: Standard Track
Expires: April 23, 2014

T. Choi
ETRI
J. Shinn
Solbox Inc.
J. Sung
KT
J. Lee
SKT
J. Koo
LGU+
Y. Bang
KAIST

October 20, 2013

CDNi Control - Initialization and Bootstrapping
draft-choi-cdni-control-init-bootstrapping-02

Abstract

This document proposes a mechanism for a CDN to initiate the interconnection across CDNs and bootstrap the other CDNi interfaces.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/lid-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

Copyright and License 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
2.	Terminology	6
3.	CDNI Initialization and Bootstrapping Control Interface	6
4.	Interface Information Elements	7
4.1.	Information Element Data Type Descriptions	7
4.1.1.	ActionType	7
4.1.2.	General	7
4.1.3.	Contract	8
4.1.4.	EndPoints	8
4.1.5.	Protocols	8
4.1.6.	Fp_cap	8
4.1.7.	Cdmd	8
4.1.8.	AbsoluteTime	8
4.1.9.	ActionStatus	8
4.2.	JSON Encoding of Information Element Data Types	8
4.2.1.	ActionType	8
4.2.2.	General	9
4.2.3.	Contracts	9
4.2.4.	End-Points	9
4.2.5.	Protocols	9
4.2.6.	Fp_cap	9
4.2.7.	Cdmd	9
4.2.8.	AbsoluteTime	10
4.2.9.	ActionStatus	10
5.	Examples	10
5.1.	Discover	10
5.2.	Negotiate	11
5.3.	Initialize	11
5.4.	Add	13
5.5.	Update	14
5.6.	Remove	15

6. Security Considerations	16
7. IANA Considerations	16
8. References	17
8.1. Normative References	17
8.2. Informative References	17
Authors' Addresses	18

1 Introduction

[I-D.ietf-cdni-control-triggers-00] specifies mechanisms for CDN interconnection control for the "High" and "Medium" priority requirements for the CDNI Control Interface identified in section 4 of [I-D.ietf-cdni-requirements].

This draft concentrates on the remaining "Low" priority requirements for the CDNI Control Interface, reproduced here for convenience:

CNTL-5 [LOW] The CDNI Control interface may allow a CDN to establish, update and terminate a CDN interconnection with another CDN whereby one CDN can act as a Downstream CDN for the other CDN (that acts as an Upstream CDN).

CNTL-6 [LOW] The CDNI Control interface may allow control of the CDNI interconnection between any two CDNs independently for each direction (i.e. For the direction where CDN1 is the Upstream CDN and CDN2 is the Downstream CDN, and for the direction where CDN2 is the Upstream CDN and CDN1 is the Downstream CDN).

CNTL-7 [LOW] The CDNI Control interface may allow bootstrapping of the Request-Routing interface. For example, this can potentially include:

- * negotiation of the Request-Routing method (e.g. DNS vs HTTP, if more than one method is specified)
- * discovery of the Request-Routing protocol endpoints
- * information necessary to establish secure communication between the Request-Routing protocol endpoints.

CNTL-8 [LOW] The CDNI Control interface may allow bootstrapping of the CDNI Metadata interface. This information could, for example, include:

- * discovery of the CDNI Metadata signaling protocol endpoints
- * information necessary to establish secure communication between the CDNI Metadata signaling protocol endpoints.

CNTL-9 [LOW] The CDNI Control interface may allow bootstrapping of the Content Acquisition interface. This could, for example, include exchange and negotiation of the Content Acquisition protocols to be used across the CDNs (e.g. HTTP, HTTPS, FTP, ATIS C2).

CNTL-10 [LOW] The CDNI Control interface may allow exchange and negotiation of delivery authorization mechanisms to be supported across the CDNs (e.g. URI signature based validation).

CNTL-11 [LOW] The CDNI Control interface may allow bootstrapping of the CDNI Logging interface. This information could, for example, include:

- * discovery of the Logging protocol endpoints
 - * information necessary to establish secure communication between the Logging protocol endpoints
 - * negotiation/definition of the log file format and set of fields to be exported through the Logging protocol, with some granularity (e.g. On a per content type basis).
 - * negotiation/definition of parameters related to transaction Logs export (e.g., export protocol, file compression, export frequency, directory).

This document consider those parts of the control interface that relate to configuration, bootstrapping or authentication of CDN Interconnect interfaces.

- o Section 2 terminology used in this document..
- o Section 3 defines the RESTful web service interface used for CDNi initiation and bootstrapping.
- o Section 4 defines information elements encapsulated in the request and response messages of the interface and their JSON encoding
- o Section 5 contains example messages.

2. Terminology

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].

3. CDNI Initialization and Bootstrapping Control Interface

In order to meet the above requirements, the control interface should support the following actions.

- o discover - used to instruct a recipient CDN to return the requested information to the originator. The recipient CDN then fill in the requested information for the discovery.
- o negotiate - used to instruct a recipient CDN to negotiate the requested information (e.g., authentication & request-routing methods, content acquisition protocol, log file format, log export method). The recipient CDN checks with originator's preference and fill in its responding preference.
- o initialize - used to instruct a recipient CDN to initialize a CDN interconnection with all other CDNs at boot time.
- o add - used to instruct a recipient CDN to add newly created footprint & capabilities(fp_cap) or content and metadata(cdmd).
- o update - used to instruct a recipient CDN to update the existing fp_cap or cdmd when there are changes.
- o remove - used to instruct a recipient CDN to remove the existing fp_cap or cdmd.

The CDNI initialization and bootstrapping control interface is a simple RESTful interface based on HTTP[RFC2616]. The attributes and actions are encapsulated in the messages.

Figure 1 shows the generic message flow used between the originating CDN and the recipient CDN for initialization and bootstrapping operations.

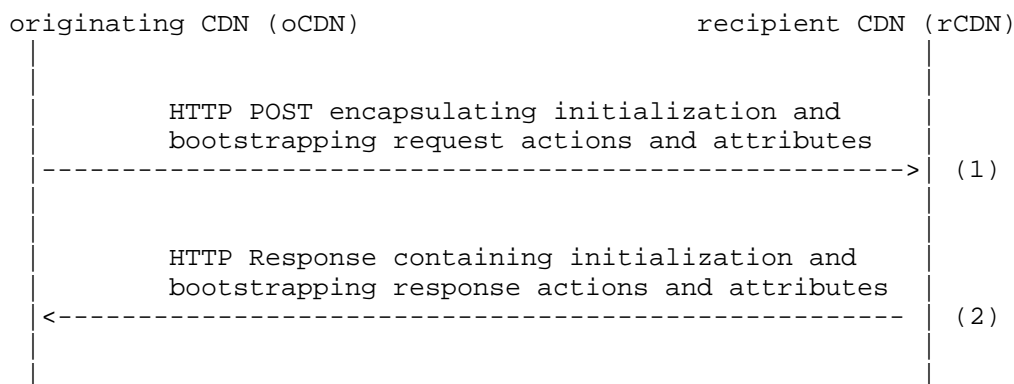


Figure 1: Basic CDNI Message Flow for Triggers

The steps in Figure 1 are:

1. originating CDN sends a HTTP POST to recipient CDN containing the action and attributes of initialization or bootstrapping.
2. recipient CDN authenticates and validates the request and if it accepts the request, performs actions requested, and responds to originating CDN with an HTTP "200" response with a message body containing the performed results and the associated resulting attributes.

4. Interface Information Elements

Information exchanged between originating CDN and recipient CDN consists of the following.

4.1. Information Element Data Type Descriptions

This section describes the data types that are used for properties of information elements.

4.1.1. ActionType

This type defines the type of action being requested, permitted actions are discover, initialize, negotiate, add, update, and remove.

4.1.2. General

This type describes a set of general information of a CDN, possible values are cdnname, cdnid, cdnhostname, cdnigw, and cdnserver.

4.1.3. Contract

This type describes a set of CDN information where interconnection contract is established. It consists of cdnname, cdnid, cdnhostname, cdnigw, ucdnflag, and dcdnflag. xcdnflag indicates whether it is capable of either uCDN or dCDN.

4.1.4. EndPoints

This type describes a set of end points of CDNi interfaces. It consists of request routing protocol, metadata signaling protocol, and logging protocol end points.

4.1.5. Protocols

This type describes protocols used by various CDNi interfaces. It consists of request routing redirection, logging, metadata signaling, and content acquisition.

4.1.6. Fp_cap

This type describes a set of footprint and capabilities information needed for bootstrapping. It consists of cdnname, iprange, delay, load, and bandwidth.

4.1.7. Cdm

This type describes a set of contents and metadata information needed for bootstrapping. It consists of hostindex, hostmetadatas, pathmetadatas.

4.1.8. AbsoluteTime

This type describes an absoluteTime. Times are expressed in seconds since the UNIX epoch.

4.1.9. ActionStatus

This type describes resulting status of the requested action. It indicates either success or failure.

4.2. JSON Encoding of Information Element Data Types

4.2.1. ActionType

Key: type

Description: One of "discover", "negotiate", "initialize", "add",

"update", or "remove".

Type: String

Mandatory: Yes

4.2.2. General

Keys: general

Description: A set of general information of a CDN.

Type: String

Mandatory: Yes

4.2.3. Contracts

Key: contracts

Description: A set of CDN information where the interconnection contract is established.

Type: String

Mandatory: Yes

4.2.4. End-Points

Keys: ed_points

Description: A set of CDNi interface protocols end-point information for bootstrapping.

Type: String

Mandatory: Yes

4.2.5. Protocols

Keys: protocols

Description: A set of CDNi protocols used for various interfaces.

Type: String

Mandatory: Yes

4.2.6. Fp_cap

Keys: fp_cap

Description: A set of footprint and capability information needed for bootstrapping.

Type: String

Mandatory: Yes

4.2.7. Ccmd

Keys: ccmd

Description: A set of content and metadata information needed for bootstrapping.

Type: String

Mandatory: Yes

4.2.8. AbsoluteTime

Keys: AbsoluteTime

Description:

Type: String

Mandatory: Yes

4.2.9. ActionStatus

Keys: ActionStatus

Description: Integer 1 or 0 to indicate the success or failure of requested actions.

Type: Integer

Mandatory: Yes

5. Examples

The following sections provide examples of different CI/B objects encoded as JSON.

No authentication is shown in the following illustrative examples, it is anticipated that authentication mechanisms will be aligned with other CDNI Interfaces as and when those mechanisms are defined.

5.1. Discover

REQUEST:

POST /bootstrap HTTP/1.1

Host: dcdn.example.com

Accept: application/vnd.cdni.control.boot.request+json

Content-Length: xx

```
{
  "ActionType": "discover",
  "cdmd.endpoints" : "http://metadata.ucdn.com/ep",
  "rr.endpoints"   : "http://rr.ucdn.com/ep",
  "log.endpoints"  : "http://log.ucdn.com/ep"
}
```

RESPONSE:

HTTP/1.1 200 OK

Date: Sat, 19 Oct 2013 14:20:08 GMT

Content-Type: application/vnd.cdni.control.boot.response+json

```
{
  "cdmd.ednpoints": "http://metadata.dcdn.com/ep",
  "rr.endpoints"  : "http://rr.dcdn.com/ep",
  "log.endpoints" : "http://log.dcdn.com/ep",
}
```



```
        "actionStatus" : 1
    }
```

5.2. Negotiate

REQUEST:

```
POST /bootstrap HTTP/1.1
Host: dcdn.example.com
Accept: application/vnd.cdni.control.boot.request+json
Content-Length: xx
{
    "ActionType": "negotiate",
    "cdmd.protocol" : "HTTP",
    "rr.protocol" : "DNS",
    "log.protocol" : "FTP"
}
```

RESPONSE:

```
HTTP/1.1 200 OK
Date: Sat, 19 October 2013 14:20:06 GMT
Content-Type: application/vnd.cdni.control.boot.response+json
{
    "cdmd.protocol" : "HTTP",
    "rr.protocol" : "DNS",
    "log.protocol" : "FTP",
    "actionStatus" : 1
}
```

5.3. Initialize

```
POST /bootstrap HTTP/1.1
Host: dcdn.example.com
Accept: application/vnd.cdni.control.boot.request+json
Content-Length: xx
{
    "ActionType": "initialize",
    "cdni.general" : [
        {
            "cdnname": "UCDN1",
            "cdnid": 100000,
            "cdnhostname": "ucdn.example.com"
        }
    ],
    "cdni.contract" : [
        {
            "cdnname": "DCDN1",
            "cdnid": 200000,

```

```

        "cdnhostname": "dcdn.example.com"
        "dcdnflag": 1
        "ucdnflag": 0
    }
],
"cdni.fp_cap" : [
    {
        "cdnname": "UCDN1",
        "iprange": 0.0.0.0/32,
        "delay": 10,
        "load": 30,
        "bandwidth": 20
    },
    {
        "cdnname": "UCDN1",
        "iprange": 211.224.204.0/24,
        "delay": 10,
        "load": 30,
        "bandwidth": 20
    }
],
"cdni.cdmd" : [
    {
        "HostIndex": [
            {
                "hosts": [
                    {
                        "host": "video.example.com",
                        "links": [
                            {
                                "rel": "host-metadata",
                                "type": "application/cdni.HostMetadata",
                                "href": "http://metadata.example.ucdn.com
/video"
                            }
                        ]
                    },
                    {
                        "host": "images.example.com",
                        "links": [
                            {
                                "rel": "host-metadata",
                                "type": "application/cdni.HostMetadata",
                                "href": "http://metadata.ucdn.example.com
/images"
                            }
                        ]
                    }
                ]
            }
        ]
    }
]

```

```
    ]
  }
}
}
```

RESPONSE:

HTTP/1.1 200 OK

Date: Sat, 19 October 2013 14:20:08 GMT

Content-Type: application/vnd.cdni.control.boot.status+json

Content-Length: xx

```
{
  "cdni.general" [ .. ],
  "cdni.contract" [ .. ],
  "cdni.fp_cap" [ .. ],
  "cdni.cdmd" [ .. ],
  "ActionStatus": 1
}
```

5.4. Add

POST /bootstrap HTTP/1.1

Host: dcdn.example.com

Accept: application/vnd.cdni.control.boot.request+json

Content-Length: xx

```
{
  "ActionType": "add",
  "cdni.fp_cap" : [
    {
      "cdnname": "UCDN1",
      "iprange": 211.224.205.0/24,
      "delay": 30,
      "load": 10,
      "bandwidth": 70
    }
  ],
  "cdni.cdmd" : [
    {
      "HostIndex": [
        {
          "hosts": [
            {
              "host": "streaming.example.com",
              "links": [
                {
                  "rel": "host-metadata",
                  "type": "application/cdni.HostMetadata",

```

```
        "href": "http://metadata.ucdn.example.com/streaming"
      }
    ]
  }
}
}
```

RESPONSE:

HTTP/1.1 200 OK

Date: Sat, 19 October 2013 14:20:08 GMT

Content-Type: application/vnd.cdni.control.boot.status+json

Content-Length: xx

```
{
  "cdni.fp_cap" [ .. ],
  "cdni.cdmd" [ .. ],
  "ActionStatus": 1
}
```

5.5. Update

POST /bootstrap HTTP/1.1

Host: dcdn.example.com

Accept: application/vnd.cdni.control.boot.request+json

Content-Length: xx

```
{
  "ActionType": "update",
  "cdni.fp_cap" : [
    {
      "cdnname": "UCDN1",
      "iprange": "211.224.205.0/24",
      "delay": 5,
      "load": 10,
      "bandwidth": 60
    }
  ],
  "cdni.cdmd" : [
    {
      "HostIndex": [
        {
          "hosts": [
            {
              "host": "streaming.example.com",
```

```
        "links": [
          {
            "rel": "host-metadata",
            "type": "application/cdni.HostMetadata",
            "href": "http://metadata.ucdn.example.com/streaming1"
          }
        ]
      }
    ]
  }
}
}
```

RESPONSE:

HTTP/1.1 200 OK

Date: Sat, 19 October 2013 14:20:08 GMT

Content-Type: application/vnd.cdni.control.trigger.status+json

Content-Length: xx

```
{
  "cdni.fp_cap" [ .. ],
  "cdni.cdmd" [ .. ],
  "ActionStatus": 1
}
```

5.6. Remove

POST /bootstrap HTTP/1.1

Host: dcdn.example.com

Accept: application/vnd.cdni.control.trigger.request+json

Content-Length: xx

```
{
  "ActionType": "remove",
  "cdni.fp_cap" : [
    {
      "cdnname": "UCDN1",
      "iprange": 211.224.205.0/24,
      "delay": 5,
      "load": 10,
      "bandwidth": 60
    }
  ],
  "cdni.cdmd" : [
    {
      "HostIndex": [
```

```
{
  "hosts": [
    {
      "host": "streaming.example.com",
      "links": [
        {
          "rel": "host-metadata",
          "type": "application/cdni.HostMetadata",
          "href": "http://metadata.ucdn.example.com/streaming1"
        }
      ]
    }
  ]
}
```

RESPONSE:

HTTP/1.1 200 OK

Date: Sat, 19 October 2013 14:20:08 GMT

Content-Type: application/vnd.cdni.control.boot.status+json

Content-Length: xx

```
{
  "cdni.fp_cap" [ .. ],
  "cdni.cdmd" [ .. ],
  "ActionStatus": 1
}
```

6. Security Considerations

The recipient CDN must ensure that each originating CDN only has access to its own Trigger Status Resources.

It is anticipated that a common authentication mechanism will be used by this and other CDNI Interconnect interfaces, the mechanism must exist but is not identified in this document.

The recipient CDN must ensure that activity triggered by originating CDN only affects metadata or content originating from that originating CDN.

7. IANA Considerations

TBD.

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2616, June 1999.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005.

8.2. Informative References

- [I-D.ietf-cdni-control-triggers-00] Murray, R., Niven-Jenkins, B., "CDN Interconnect Triggers", draft-murray-cdni-triggers-01 (work in progress), April, 2013.
- [I-D.cjlmw-cdni-metadata] Niven-Jenkins, B., Murray, R., Watson, G., Caulfield, M., and K. Leung, "CDN Interconnect Metadata", draft-ietf-cdni-metadata-01 (work in progress), February 2013.
- [I-D.ietf-cdni-framework] Peterson, L. and B. Davie, "Framework for CDN Interconnection", draft-ietf-cdni-framework-03 (work in progress), February 2013.
- [RFC6707] Niven-Jenkins, B., Faucheur, F., and N. Bitar, "Content Distribution Network Interconnection (CDNI) Problem Statement", September 2012.
- [I-D.ietf-cdni-requirements] Leung, K. and Y. Lee, "Content Distribution Network Interconnection (CDNI) Requirements", draft-ietf-cdni-requirements-09 (work in progress), July 2013.
- [RFC4287] Nottingham, M., Ed. and R. Sayre, Ed., "The Atom Syndication Format", RFC 4287, December 2005.

[XML-BASE]

Marsh, J., Ed. and R. Tobin, Ed., "XML Base (Second Edition) - <http://www.w3.org/TR/xmlbase/>", January 2009.

Authors' Addresses

Taesang Choi
ETRI
161 Gajong-Dong
Yusong-Gu, Daejeon
Republic of Korea

Phone: +82-42-860-5628
Email: choits@etri.re.kr

John Dongho Shinn
Solbox Inc.
7F, Haesung Bldg. 747-2 Yeoksam-Dong
Kangnam-Gu, Seoul
Republic of Korea

Phone: +82-10-3005-4785
Email: eastsky@solbox.com

Jonggyu Sung
KT Infra Laboratory
463-1, Jeonmin-dong
Yuseong-gu, Daejeon
Republic of Korea

Phone: 82-10-3096-8237
Email: jonggyu.sung@kt.com

Jongmin Lee
SK Telecom
11, Euljiro-2ga
Jung-gu, Seoul
Republic of Korea

Phone: 82-10-9429-6260
Email: jminlee@sk.com

Ja-Ryeong Koo
LG U plus Corporation
Namdaemunro 5-ga
Jung-gu, Seoul
Republic of Korea

Phone: 82-10-8080-6115
Email: wjbkoo@lguplus.co.kr

Yong Hwan Bang
KAIST
8F, N29 Bldg. 373-1 Gusung-Dong
Yousong-Gu, Daejeon
Republic of Korea

Phone: +82-42-350-7516
Email: yjhvyp@gmail.com

Network Working Group
Internet-Draft
Intended Status: Standard Track
Expires: April 20, 2014

T. Choi
ETRI
J. Shinn
Solbox Inc.
J. Sung
KT
J. Lee
SKT
J. Koo
LGU+
Y. Bang
KAIST

October 17, 2013

Request Interface for Iterative Request Routing Redirection
draft-choi-cdni-req-intf-01

Abstract

This document specifies the request interface between a CDN end-user and an upstream CDN or downstream CDN to request CDN request routing redirection. It specifies the CDNI Request Interface information elements and the actual protocol for exchanging them.

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 December 29, 2012.

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	3
1.1. Terminology	3
2. Request Interface Function Overview	3
3. Request Interface Information Elements	3
4. Request Procedures	4
4.1. HTTP-based Iterative Request Routing Redirection with Loop Prevention	4
4.2. DNS-based Iterative Request Routing Redirection with Loop Prevention	9
5 Security Considerations	11
6 IANA Considerations	12
7. References	12
7.1. Normative References	12
7.2. Informative References	12
Authors' Addresses	12

1. Introduction

According to the CDNI requirements [I-D.ietf-cdni-requirements] and the CDNI Framework [I-D.ietf-cdni-framework], the need for interconnected CDNs to be able to support iterative and recursive CDNI request routing and implement an access control mechanism that enforces the CSP's distribution policy is described. Recursive CDNI request routing is specified by request routing redirection interface [I-D.ietf-cdni-redirection]. URI signing is specified in URI signing [I-D.leung-cdni-uri-signing].

Meanwhile, both iterative request routing redirection and URI signing require interactions between an end-user and uCDN or dCDN. This implies the need for an interface to support such interactions.

To meet the needs, this document specifies an interface called "Request Interface" which resides between end-users and CDN providers. It describes information elements and transport protocols for the interface.

1.1. Terminology

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].

2. Request Interface Function Overview

Request interface operates between a user agent and CDN provider (either uCDN or dCDN). Since one end of this interface is user application such as web browser, it is desirable not to invent a new protocol for this function. Instead, it uses well-known protocols used by the existing CDN applications such as HTTP and DNS.

The main function of this interface is to support the iterative request routing redirection and URI signing. For the iterative request routing redirection, it also supports loop prevention mechanism.

3. Request Interface Information Elements

To meet the requirements of request routing redirection, we define the term "CDN-Provider-ID". It uniquely identifies each CDN provider during the course of request routing redirection. It consists of "CDN provider name" and "MaxNumRedHops". A pair of an AS number and an additional qualifier is used for CDN provider name. Since more than one CDN providers can belong to the same AS, an additional qualifier is used to guarantee the uniqueness. MaxNumRedHops

represents a maximum allowed redirections. The value is decreased once every redirection occurs until it reaches 0. To avoid its usage abuse (e.g., end user or CDN operator can set huge number like 100 or above), a reasonable upper bound has to be agreed among CDN providers. Security aspect of it is for further study.

A few examples of the CDN provider names are 100:0 and 200:1. The former means that a CDN provider belong to AS 100 and it is the only CDN provider within that AS. The latter represents the first CDN provider in the AS 200. There are other CDN providers in the same AS.

One example of CDN-Provider-ID is "CDN-Provider-Name=100:0 & MaxNumRedHops=10", which means that a CDN provider that belong to AS number 100 and it is the only CDN provider and a maximum allowed redirection is 10. An example how a list of CDN-Provider-IDs can be carried in the URI query string when a certain cascaded request routing redirection occurs is the following. We assume that redirection is cascaded three times: uCDN -> dCDN1 -> dCDN2. dCDN1, then, carries the following URL, "http://cdn.csp.com?uCDN-Provider-ID=100:0&dCDN1-Provider-ID=200:1&MaxNumRedHops=9". Note that MaxNumRedHops carries the latest number instead of adding in every CDN-Provider-ID to save the space in URI query string.

It is applicable for both HTTP-based and DNS redirections. For HTTP-based redirection, we define a HTTP request routing redirection header "CDN-Provider-ID". For each step of redirection, it is attached to the beginning of the provider domain URL. For example, uCDN initiates a redirection with its URL, http://100:0:10.cdn.csp.com. dCDN further attaches its own CDN-Provider-ID in the front when another level of redirection is required. For DNS-based redirection, the CDN-Provider-ID can be attached in the DNS CNAME.

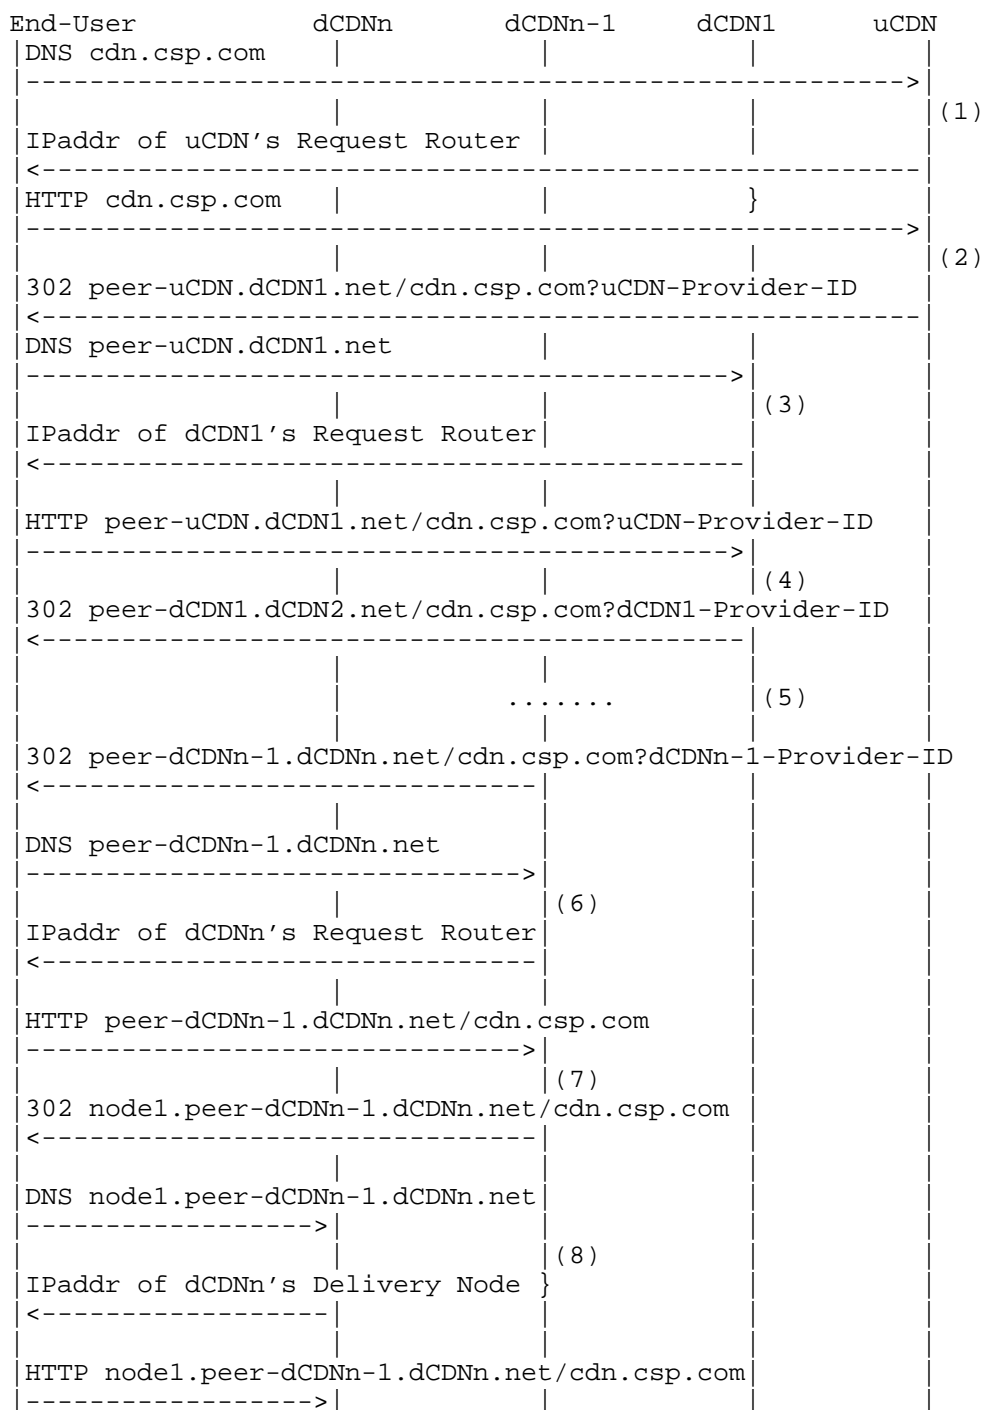
For URI signing, detailed format of information elements carried over this interface is defined in [I-D.ietf-leung-cdni-uri-signing-02]

4. Request Procedures

This section specifies two request procedures: HTTP-based and DNS-based iterative request routing redirection with loop prevention.

4.1. HTTP-based Iterative Request Routing Redirection with Loop Prevention

In this section, we describe an iterative procedure of HTTP-based request routing redirection with loop prevention.



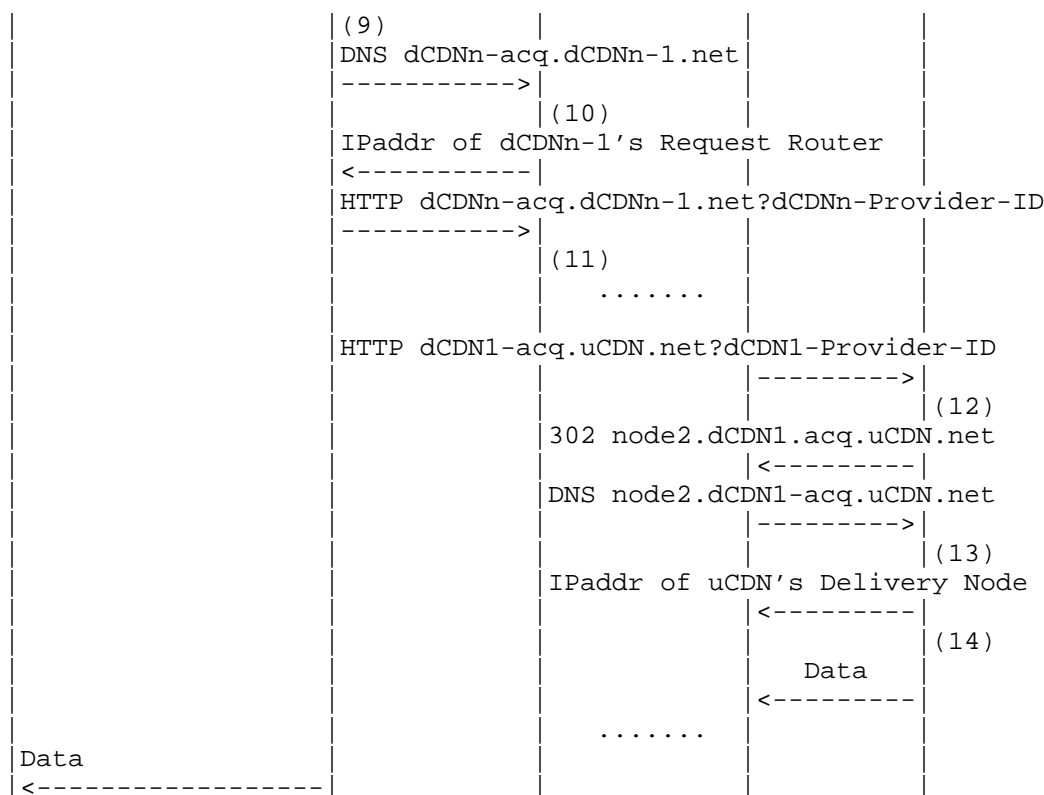


Figure 1: HTTP-based request routing redirection iterative procedure

The steps illustrated in the figure are as follows:

1. A DNS resolver for uCDN provider processes the DNS request for its customer based on CDN-domain `cdn.csp.com`. It returns the IP address of a request router in uCDN provider.
2. A Request Router for uCDN provider processes the HTTP request and recognizes that the end-user is best served by dCDN1. So it returns a 302 redirect message for a new URL constructed by "stacking" dCDN1's distinguished CDN-domain (`peer-uCDN.dCDN1.net`) on the front of the original URL. It also adds uCDN's CDN-Provider-ID in the URI query string of the HTTP request message. (e.g., `uCDN-Provider-ID=100:0 & MaxNumRedHops=10`). This information is not processed by the customer but conveyed in the HTTP message without any modification of the step 4. The details on how it is used for loop prevention is described in the step 4.

3. The end-user does a DNS lookup using dCDN1's distinguished CDN-domain (peer-uCDN.dCDN1.net). dCDN1's DNS resolver returns the IP address of a request router for dCDN1.
4. The request router for dCDN1 processes the HTTP request. There are two options: redirect further to another dCDN (i.e., cascading the request) or process it by itself. In either cases, it performs loop prevention step first. It checks a list of CDN-provider-IDs in the URI query string: it contains a list of CDN providers which requested redirections so far. If either it contains own CDN provider name or MaxNumRedHops becomes 0, it means that the redirection loop has occurred or the number of redirection hops has reached the maximum. Once loop is detected, details on the next steps is described in the section 3. If it is loop free, it either redirects further or processes based on the local policy. For the former, it selects another dCDN provider and sends an HTTP redirect message with its own CDN-Provider-ID included in its URI query string (e.g., uCDN-Provider-ID=100:0 & dCDN1-Provider-ID=200:1 & MaxNumRedHops=9) attached. For the latter, it selects a suitable delivery node to serve the end-user request, and returns a 302 redirect message for a new URL constructed by replacing the hostname by a subdomain of the dCDN1's distinguished CDN-domain that points to the selected delivery node. Then it goes to the step 6.
5. If further redirection is decided, it repeats steps 2 - 4 until it either selects dCDN provider to serve the request or MaxNumRedHops expires. If the former occurs, it resumes the step 6. If the latter occurs, it follows the processes described in the section 3.
6. Assuming that dCDNn is selected as a serving dCDN provider, the end-user does a DNS lookup using dCDNn's distinguished CDN-domain (peer-dCDNn-1.dCDNn.net). dCDNn-1's DNS resolver returns the IP address of a request router for dCDNn.
7. The request router for dCDN1 processes the HTTP request and selects a suitable delivery node to serve the end-user request, and returns a 302 redirect message for a new URL constructed by replacing the hostname by a subdomain of the dCDNn's distinguished CDN-domain that points to the selected delivery node.
8. The end-user does a DNS lookup using dCDNn's delivery node subdomain (node1.peer-dCDNn-1.dCDNn.net). dCDNn's DNS resolver returns the IP address of the delivery node.

9. The end-user requests the content from dCDNn's delivery node. In the case of a cache hit, steps 10 ~ 14 below do not happen, and the content data is directly returned by the delivery node to the end-user. In the case of a cache miss, the content needs to be acquired by dCDN from either parent dCDN or uCDN (not the CSP). The distinguished CDN-domain peer-dCDNn-1.dCDNn.net indicates that this content is to be acquired from dCDNn-1; stripping the CDN-domain reveals the original CDN-domain cdn.csp.com and dCDNn may verify that this CDN-domain belongs to a known peer (so as to avoid being tricked into serving as an open proxy). It then does a DNS request for an inter-CDN acquisition CDN-domain as agreed above (in this case, dCDNn-acq.dCDNn-1.net). This process repeats recursively until it finds a CDN provider that can serve the requested content.
10. dCDNn-1's DNS resolver processes the DNS request and returns the IP address of a request router in dCDNn-1.
11. The request router for dCDNn-1 processes the HTTP request from dCDNn's delivery node. dCDNn-1 request router recognizes that the request is from a peer CDN rather than an end-user because of the dedicated inter-CDN acquisition domain (dCDNn-acq.dCDNn-1.net). It also performs loop prevention process as described in step 4 based on the provided CDN-Provider-ID (e.g., uCDN-Provider-ID=100:0 & dCDN1-Provider-ID=200:1 & ... & dCDNn-Provider-ID=1000:0 & MaxNumRedHops=1). Depending on the number of levels of redirection and availability of contents, the same process repeats until either content serving CDN provider is found or MaxNumRedHops expires.
12. Assuming that all intermediate dCDNs also have a cache miss, The request router for uCDN selects a suitable delivery node to serve the inter-CDN acquisition request and returns a 302 redirect message for a new URL constructed by replacing the hostname by a subdomain of the uCDN's distinguished inter-CDN acquisition domain that points to the selected delivery node.
13. uCDN DNS resolver processes the DNS request and returns the IP address of the delivery node in uCDN.
14. uCDN serves content for the requested CDN-domain to dCDN and finally to end-user. Although not shown, it is at this point that uCDN processes the rest of the URL: it extracts information identifying the origin server, validates that this server has been registered, and determines the content provider that owns the origin server. It may also perform its own content acquisition steps if needed before returning the content to dCDN.

4.2. DNS-based Iterative Request Routing Redirection with Loop Prevention

In this section, we describe an iterative procedure of DNS-based request routing redirection with loop prevention.

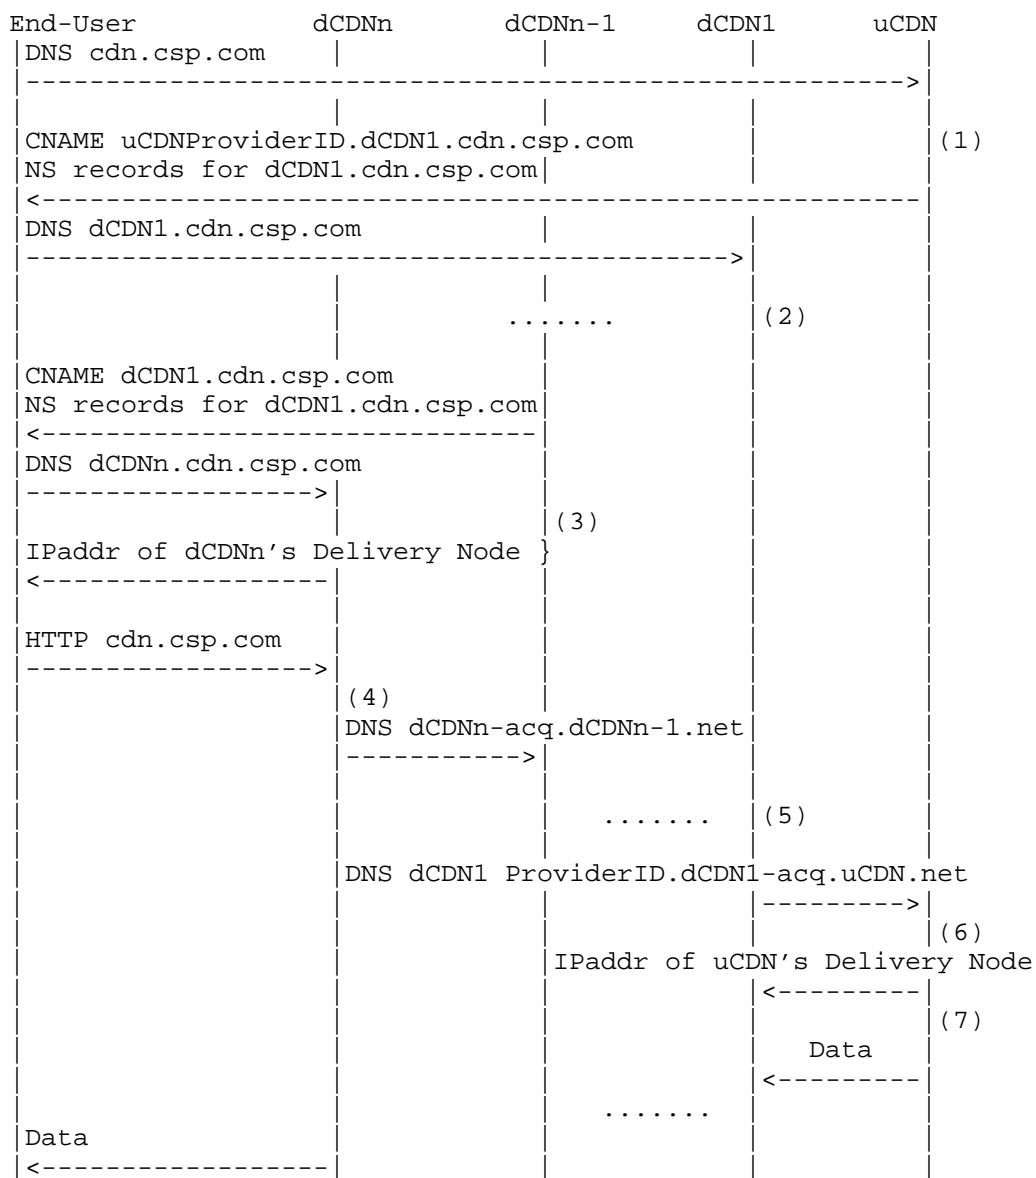


Figure 2: DNS-based request routing redirection iterative procedure

The steps illustrated in the figure are as follows:

1. Request Router for uCDN provider processes the DNS request for CDN- domain `cdn.csp.com` and recognizes that the end-user is best

served by another CDN. (This may depend on the IP address of the user's local DNS resolver, or other information discussed below.) The Request Router returns a DNS CNAME response by "stacking" the distinguished identifier for dCDN1 and uCDN's CDN-Provider-ID (e.g., 100:0.10) onto the original CDN-domain (e.g., dCDN1.cdn.csp.com), plus an NS record that maps dCDN1.cdn.csp.com to dCDN1's Request Router.

2. The end-user does a DNS lookup using the modified CDN-domain (i.e., dCDN1.cdn.csp.com). dCDN1 Request Router processes the request and decides to serve the request or redirect further to another CDN provider. It also checks redirection loop. This process iterates until either serving dCDN is selected or MaxNumRedHops expires. In this case, dCDNn is selected as a serving dCDN. If the former occurs, it proceeds to step 3. If the latter occurs, it follows the processes described in the section 3.
3. The end-user does a DNS lookup using the modified CDN-domain (i.e., dCDN1.cdn.csp.com). This causes dCDNn's request router returns an IP address of a suitable delivery node.
4. The end-user requests the content from dCDNn's delivery node. In the case of a cache hit, steps 5 ~ 7 do not happen, and the content data is directly returned by the delivery node to the end-user. In the case of a cache miss, the content needs to be acquired by dCDNn from either parent dCDN or uCDN (not the CSP). It also performs loop prevention process as described in the step 2 based on the provided CDN-Provider-ID (e.g., 100:0.200:1.....900:0.1)
5. Depending on the number of levels of redirection and availability of contents, the same process repeats until either content serving CDN provider is found or MaxNumRedHops expires.
6. Assuming that all intermediate dCDNs also miss cache, uCDN is selected as a content delivery CDN provider. Thus, the request router for uCDN selects a suitable delivery node to serve the inter-CDN acquisition request and returns IP address of the suitable uCDN delivery node.
7. uCDN serves content to dCDN1 and further down to end-user. Although not shown, it is at this point that uCDN processes the rest of the URL: it extracts information identifying the origin server, validates that this server has been registered, and determines the content provider that owns the origin server.

5 Security Considerations

TBD.

6 IANA Considerations

TBD.

7. References

7.1. Normative References

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

[I-D.ietf-cdni-framework] Peterson, L. and B. Davie, "Framework for CDN Interconnection", February 2013.

[I-D.ietf-cdni-requirements] Leung, K. and Y. Lee, "Content Distribution Network Interconnection (CDNI) Requirements", July 2013.

[I-D.ietf-cdni-redirection] Wang, H. and Niven-Jenkins, B., "Request Routing Redirection Interface for Content Distribution Network Interconnection", April 2013.

[I-D.leung-cdni-uri-signing] Leung, K. and Le Faucheur, F., "URI Signing for Content Distribution Network Interconnection", May 2013.

7.2. Informative References

TBD.

Authors' Addresses

Taesang Choi
ETRI
161 Gajong-Dong
Yusong-Gu, Daejeon
Republic of Korea

Phone: +82-42-860-5628
Email: choits@etri.re.kr

John Dongho Shinn
Solbox Inc.

7F, Haesung Bldg. 747-2 Yeoksam-Dong
Kangnam-Gu, Seoul
Republic of Korea

Phone: +82-10-3005-4785
Email: eastsky@solbox.com

Jonggyu Sung
KT Infra Laboratory
463-1, Jeonmin-dong
Yuseong-gu, Daejeon
Republic of Korea

Phone: 82-10-3096-8237
Email: jonggyu.sung@kt.com

Jongmin Lee
SK Telecom
11, Euljiro-2ga
Jung-gu, Seoul
Republic of Korea

Phone: 82-10-9429-6260
Email: jminlee@sk.com

Ja-Ryeong Koo
LG U plus Corporation
Namdaemunro 5-ga
Jung-gu, Seoul
Republic of Korea

Phone: 82-10-8080-6115
Email: wjbkoo@lguplus.co.kr

Yong Hwan Bang
KAIST
8F, N29 Bldg. 373-1 Gusung-Dong
Yousong-Gu, Daejeon
Republic of Korea

Phone: +82-42-350-7516
Email: yj hvyp@gmail.com

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 24, 2014

R. Murray
B. Niven-Jenkins
Velocix (Alcatel-Lucent)
October 21, 2013

CDNI Control Interface / Triggers
draft-ietf-cdni-control-triggers-01

Abstract

This document describes the part of the CDN Interconnect Control Interface that allows a CDN to trigger activity in an interconnected CDN that is configured to deliver content on its behalf. The upstream CDN can use this mechanism to request that the downstream CDN pre-positions metadata or content, or that it re-validate or purge metadata or content. The upstream CDN can monitor the status of activity that it has triggered in the downstream CDN.

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 April 24, 2014.

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	5
2. Model for CDNI Triggers	5
2.1. Timing of Triggered Activity	7
2.2. Trigger Results	7
3. Collections of Trigger Status Resources	7
4. CDNI Trigger interface	8
4.1. Creating Triggers	10
4.2. Checking Status	11
4.2.1. Polling Trigger Status Resource collections	11
4.2.2. Polling Trigger Status Resources	11
4.3. Deleting Triggers	11
4.4. Expiry of Trigger Status Resources	12
4.5. Error Handling	12
5. Properties of Triggers	13
5.1. Properties of Trigger Requests	13
5.1.1. Content URLs	14
5.2. Properties of Trigger Status Resources	14
5.3. Properties of ErrorDesc	15
5.4. Properties of Trigger Collections	16
5.5. Trigger Resource Simple Data Type Descriptions	16
5.5.1. TriggerType	16
5.5.2. TriggerStatus	16
5.5.3. URLs	17
5.5.4. AbsoluteTime	17
5.5.5. ErrorCode	17
6. JSON Encoding of Objects	17
6.1. JSON Encoding of Embedded Types	18
6.1.1. TriggerType	18
6.1.2. TriggerStatus	18
6.1.3. PatternMatch	18
6.1.4. ErrorDesc	19
6.1.5. ErrorCode	19
6.1.6. Relationship	20
6.2. MIME Media Types	20
7. Examples	20

7.1. Creating Triggers	21
7.1.1. Preposition	21
7.1.2. Invalidate	22
7.2. Examining Trigger Status	24
7.2.1. Collection of All Triggers	24
7.2.2. Filtered Collections of Triggers	25
7.2.3. Trigger Status Resources	26
7.2.4. Polling for Change	28
7.2.5. Cancelling or Removing a Trigger	31
7.2.6. Error Reporting	33
8. IANA Considerations	34
9. Security Considerations	34
10. Acknowledgements	34
11. References	34
11.1. Normative References	34
11.2. Informative References	34
Authors' Addresses	35

1. Introduction

[RFC6707] introduces the Problem scope for CDN Interconnection (CDNI) and lists the four categories of interfaces that may be used to compose a CDNI solution (Control, Metadata, Request Routing, Logging).

[I-D.ietf-cdni-framework] expands on the information provided in [RFC6707] and describes each of the interfaces and the relationships between them in more detail.

This document describes the "CI/T" interface, "CDNI Control interface / Triggers". It does not consider those parts of the control interface that relate to configuration, bootstrapping or authentication of CDN Interconnect interfaces. Requirements for CI/T are the "High" and "Medium" priority requirements for the CI identified in section 4 of [I-D.ietf-cdni-requirements], reproduced here for convenience:

CI-1 [HIGH] The CDNI Control interface shall allow the Upstream CDN to request that the Downstream CDN, including downstream cascaded CDNs, delete an object or set of objects and/or its CDNI metadata from the CDN surrogates and any storage. Only the object(s) and CDNI metadata that pertain to the requesting Upstream CDN are allowed to be purged.

CI-2 [MED] The CDNI Control interface should allow for multiple content items identified by a Content Collection ID to be purged using a single Content Purge action.

CI-3 [MED] The CDNI Control interface should allow the Upstream CDN to request that the Downstream CDN, including downstream cascaded CDNs, mark an object or set of objects and/or its CDNI metadata as "stale" and revalidate them before they are delivered again.

CI-4 [HIGH] The CDNI Control interface shall allow the Downstream CDN to report on the completion of these actions (by itself, and including downstream cascaded CDNs, in a manner appropriate for the action (e.g. synchronously or asynchronously). The confirmation receipt should include a success or failure indication. The failure indication along with the reason are used if the Downstream CDN cannot delete the content in its storage.

CI-5 [MED] The CDNI Control interface should support initiation and control by the Upstream CDN of pre-positioned CDNI metadata acquisition by the Downstream CDN.

CI-6 [MED] The CDNI Control interface should support initiation and control by the Upstream CDN of pre-positioned content acquisition by the Downstream CDN.

- o Section 2 outlines the model for the CI/T Interface at a high level.
- o Section 3 describes collections of Trigger Resources.
- o Section 4 defines the RESTful web service provided by dCDN.
- o Section 5 lists properties of Trigger Requests and Status Resources.
- o Section 6 defines a JSON encoding for Trigger Requests and Status Resources.
- o Section 7 contains example messages.

1.1. Terminology

This document reuses the terminology defined in [RFC6707].

2. Model for CDNI Triggers

A trigger, sent from uCDN to dCDN, is a request for dCDN to do some work relating to data originating from uCDN.

The trigger may request action on either metadata or content, the following actions can be requested:

- o preposition - used to instruct dCDN to fetch metadata from uCDN, or content from any origin including uCDN.
- o invalidate - used to instruct dCDN to revalidate specific metadata or content before re-using it.
- o purge - used to instruct dCDN to delete specific metadata or content.

The CI/T interface is a RESTful web service offered by dCDN. It allows creation and deletion of triggers, and tracking of the triggered activity. When dCDN accepts a trigger it creates a resource describing status of the triggered activity, a Trigger Status Resource. The uCDN may poll Trigger Status Resources to monitor progress.

Requests to invalidate and purge metadata or content apply to all variants of that data with a given URI.

The dCDN maintains a collection of Trigger Status Resources for each uCDN, each uCDN only has access to its own collection and the location of that collection is shared when CDN interconnection is established.

To trigger activity in dCDN, uCDN will POST to the collection of Trigger Status Resources. If dCDN accepts the trigger, it creates a new Trigger Status Resource and returns its location to uCDN. To

monitor progress, uCDN may GET the Trigger Status Resource. To cancel a trigger, or remove a trigger from the collection once its activity has been completed, uCDN may DELETE the Trigger Status Resource.

In addition to the collection of all Trigger Status Resources for uCDN, uCDN shall have access to filtered views of that collection. These filtered views are defined in Section 3 and include collections of active and completed triggers. These collections provide a mechanism for polling the status of multiple jobs.

Figure 1 is an example showing the basic message flow used by the uCDN to trigger activity in dCDN, and for uCDN to discover the status of that activity. Only successful triggering is shown. Examples of the messages are given in Section 7.

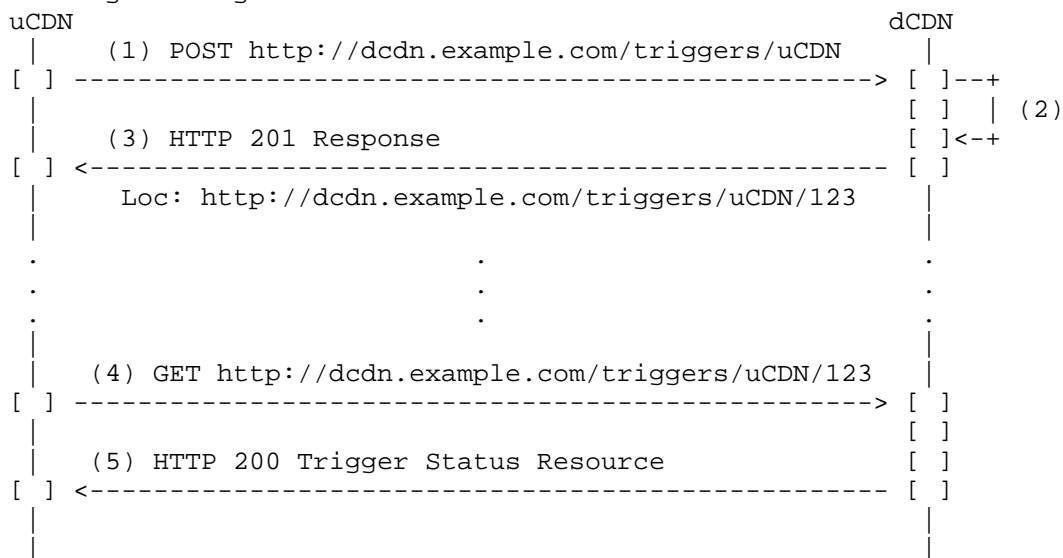


Figure 1: Basic CDNI Message Flow for Triggers

The steps in Figure 1 are:

1. uCDN triggers action in dCDN by posting to a collection of Trigger Status Resources, "http://dcdn.example.com/triggers/uCDN". The URL of this was given to uCDN when the trigger interface was established.
2. dCDN authenticates the request, validates the trigger and if it accepts the request, creates a new Trigger Status Resource.
3. dCDN responds to uCDN with an HTTP 201 response status, and the location of the Trigger Status Resource.

4. uCDN may repeatedly poll the Trigger Status Resource in dCDN.
5. dCDN responds with the Trigger Status Resource, describing progress or results of the triggered activity.

The remainder of this document describes the messages, Trigger Status Resources, and collections of Trigger Status Resources in more detail.

2.1. Timing of Triggered Activity

Timing of triggered activity is under dCDN control, including its start-time and pacing of the activity in the network.

Invalidate and purge triggers MUST be applied to all data acquired before the trigger was created in dCDN. The dCDN MAY apply the triggers to data acquired after trigger creation.

If uCDN wishes to invalidate or purge content, then immediately preposition replacement content at the same URLs, it must ensure the dCDN has completed the invalidate/purge before initiating the prepositioning. If it fails to do that and the requests overlap, and dCDN passes the triggers on to a further dCDN in a cascade, that CDN may preposition content that has not yet been invalidated/purged in its uCDN.

2.2. Trigger Results

Each Trigger Request may operate on multiple data items. The trigger shall be reported as "complete" only if all actions can be completed successfully, otherwise it shall be reported as "failed". The reasons for failure and URLs or Patterns affected shall be enumerated in the Trigger Status Resource. For more detail, see section Section 4.5.

If a dCDN is also acting as uCDN in a cascade, it MUST forward triggers to any downstream CDNs that may have data affected by the trigger. The trigger MUST NOT be reported as complete in a CDN until it is complete in all of its downstream CDNs. A trigger MAY be reported as failed as soon as it fails in a CDN or in any of its downstream CDNs.

3. Collections of Trigger Status Resources

As described in Section 2, Trigger Status Resources exist in dCDN to report the status of activity triggered by each uCDN.

A collection of Trigger Status Resources is a resource that contains

a reference to each Trigger Status Resource in that collection.

To trigger activity in dCDN, uCDN creates a new Trigger Status Resource by posting to dCDN's collection of uCDN's Trigger Status Resources. The URL of each Trigger Status Resource is generated by the dCDN when it accepts the trigger, and returned to uCDN. This immediately enables uCDN to check the status of that trigger.

The dCDN must present a different set of Trigger Status Resources to each interconnected uCDN, only Trigger Status Resources belonging to a uCDN shall be visible to it. The dCDN may, for example, achieve this by offering different collection URLs to uCDNs, or by filtering the response based on the client uCDN.

The dCDN resource representing the collection of all uCDN's Trigger Status Resources is accessible to uCDN. This collection lists all uCDN triggers that have been accepted by dCDN, and have not yet been deleted by uCDN or expired and removed by dCDN.

In order to allow uCDN to check status of multiple jobs in a single request, dCDN shall also maintain collections representing filtered views of the collection of all Trigger Status Resources. The filtered collections are:

- o Pending - Trigger Status Resources for triggers that have been accepted, but not yet acted upon.
- o Active - Trigger Status Resources for triggered activity that is currently being processed in dCDN.
- o Complete - Trigger Status Resources representing activity that completed successfully.
- o Failed - Trigger Status Resources representing activity that failed.

4. CDNI Trigger interface

This section describes an interface to enable an upstream CDN to trigger defined activities in a downstream CDN. The interface is intended to be independent of the set of activities defined now, or that may be defined in future.

CI/T is built on the principles of RESTful web services. Requests are made over HTTP, and the HTTP Method defines the operation the request would like to perform. The corresponding HTTP Response returns the status of the operation in the HTTP Status Code and returns the current representation of the resource (if appropriate) in the Response Body. HTTP Responses from servers implementing CI/T that contain a response body SHOULD include an ETag to enable validation of cached versions of returned resources.

Servers implementing CI/T MUST support the HTTP GET, HEAD, POST and DELETE methods. The only representation specified in this document is JSON.

Trigger Requests are POSTed to a URI in dCDN. If the trigger is accepted by dCDN, it creates a Trigger Status Resource and returns its URI to dCDN in an HTTP 201 response. The triggered activity can then be monitored by uCDN using that resource and the collections described in Section 3.

The URI that Trigger Requests should be POSTed to needs to be either discovered by or configured in the upstream CDN. Performing a GET on that URI retrieves a collection of the URIs of all Trigger Status Resources. The URI of each Trigger Status Resource is also returned to uCDN when it is created. This means all Trigger Status Resources can be discovered, so CI/T servers are free to assign whatever structure they desire to the URIs for CI/T resources. CI/T clients MUST NOT make any assumptions regarding the structure of CI/T URIs or the mapping between CI/T objects and their associated URIs. Therefore any URIs present in the examples below are purely illustrative and are not intended to impose a definitive structure on CI/T interface implementations.

The CI/T interface builds on top of HTTP, so CI/T servers may make use of any HTTP feature when implementing the CI/T interface. For example, a CI/T server may make use of HTTP's caching mechanisms to indicate that the returned response/representation has not been modified since it was last returned, reducing the processing needed to determine whether the status of triggered activity has changed.

This specification is neutral with regard to the transport below the HTTP layer.

[Editor's note: It is anticipated that decisions on use of HTTPS for other CDNI interfaces will be adopted for Triggers.]

Discovery of the CI/T Interface is outside the scope of this document. It is anticipated that a common mechanism for discovery of all CDNI interfaces will be defined.

The dCDN must ensure that activity triggered by uCDN only affects metadata or content originating from that uCDN. Since only one CDN can be authoritative for a given item of metadata or content, this requirement means there cannot be any "loops" in trigger requests between CDNs.

4.1. Creating Triggers

To create a new trigger, uCDN makes an HTTP POST to the unfiltered collection of its triggers. The request body of that POST is a Trigger Request.

dCDN validates and authenticates that request, if it is malformed or uCDN does not have sufficient access rights it MAY reject the request immediately. In this case, it SHALL respond with an appropriate 4xx HTTP error code and no resource shall be created on dCDN.

If the request is accepted, uCDN SHALL create a new Trigger Status Resource. The HTTP response to dCDN SHALL have status code 201 and the URI of the Trigger Status Resource in the Location header field. The HTTP response MAY include the content of the newly created Trigger Status Resource, this is recommended particularly in cases where the trigger has completed immediately.

Once a Trigger Status Resource has been created dCDN MUST NOT re-use its location, even after that resource has been removed through deletion or expiry.

The "request" property of the Trigger Status Resource SHALL contain the information posted in the body of the Trigger Request. Note that this need not be a byte-for-byte copy. For example, in the JSON representation the dCDN may re-serialise the information differently.

If the trigger is queued by dCDN for later action, the "status" property of the Trigger Status Resource SHALL be "pending". Once trigger processing has started the "status" SHALL be "active".

A trigger may result in no activity in dCDN if, for example, it is an invalidate or purge request for data the dCDN has not acquired, or a prepopulate request for data it has already acquired. In this case, the "status" of the Trigger Status Resource shall be "complete" and the Trigger Status Resource shall be added to the dCDN collection of Complete Triggers.

If dCDN is not able to track triggered activity, it MAY indicate that it has undertaken to complete the activity but will not report completion or any further errors. To do this, it must set the trigger status to "complete", with an estimated completion time in the future ("etime" greater than "mtime").

Once created, Trigger Status Resources may be deleted by uCDN but not modified. The dCDN MUST reject PUT and POST requests from uCDN to Trigger Status Resources using HTTP status code 403.

4.2. Checking Status

The uCDN has two ways to check progress of activity it has triggered in dCDN, described in the following sections.

Because the triggers protocol is based on HTTP, Entity Tags may be used by the uCDN as cache validators, as defined in section 3.11 of [RFC2616], to check for change in status of a resource or collection of resources without re-fetching the whole resource or collection.

The dCDN should use the cache control headers for responses to GETs for Trigger Status Resources and Collections to indicate the frequency at which it recommends uCDN should poll for change.

4.2.1. Polling Trigger Status Resource collections

uCDN can fetch the collection of its Trigger Status Resources, or filtered views of that collection.

This makes it possible to poll status of all triggered activity in a single request. If dCDN moves a Trigger Status Resource from the Active to the Completed collection, uCDN may chose to fetch the result of that activity.

When polling in this way, uCDN may choose to use HTTP Entity Tags to monitor for change, rather than repeatedly fetching the whole collection.

4.2.2. Polling Trigger Status Resources

uCDN has a reference (URI provided by the dCDN) for each Trigger Status Resource it has created, it may fetch that resource at any time.

This may be used to retrieve progress information, and to fetch the result of triggered activity.

4.3. Deleting Triggers

The uCDN MAY delete Trigger Status Resources at any time, using the HTTP DELETE method.

Once deleted, the references to a Trigger Status Resource MUST be removed from all Trigger Status Resource collections. Subsequent requests for the resource shall be handled as required by HTTP, and so will receive responses with status 404 or 410.

If a "pending" Trigger Status Resource is deleted, dCDN SHOULD NOT

start processing of that activity. Deleting a "pending" trigger does not however guarantee that it is not started because, once it has triggered activity, uCDN cannot control the timing of that activity. Processing may, for example, start after the DELETE is sent by uCDN and before the DELETE is processed by dCDN.

If an "active" Trigger Status Resource is deleted, dCDN MAY stop processing the triggered activity. However, as with deletion of a "pending" trigger, dCDN does not guarantee this.

Deletion of a "complete" or "failed" Trigger Status Resource requires no processing in dCDN other than deletion of the resource.

4.4. Expiry of Trigger Status Resources

The dCDN MAY choose to automatically delete Trigger Status Resources some time after they become completed or failed. In this case, dCDN will remove the resource and respond to subsequent requests for it with HTTP status 404 or 410.

If dCDN performs this housekeeping, it MUST have reported the length of time after which completed Trigger Status Resources become stale via a property of the collection of all Trigger Status Resources. It is recommended that Trigger Status Resources are automatically deleted 24 hours after they become completed or failed.

To ensure it has access to the status of its completed and failed triggers, it is recommended that uCDN's polling interval is half the time after which records for completed activity will be considered stale.

4.5. Error Handling

A CI/T server may reject a trigger request using HTTP status codes, for example 400 if the request is malformed or 401 if the client does not have permission to create triggers or it is trying to act on another CDN's data.

If any part of the trigger request fails the trigger shall be reported as "failed" once its activity is complete, or if no further errors will be reported. The "errors" property in the Trigger Status Resource will be used to enumerate which actions failed and the reasons for failure, and may be present while the trigger is still "pending" or "active" if the trigger is still running for some URLs or Patterns in the trigger request.

Once a request has been accepted, processing errors are reported in the Trigger Status Resource using a list of "ErrorDesc". Each

ErrorDesc is used to report errors against one or more of the URLs or Patterns in the trigger request.

If a surrogate affected by a trigger is offline in dCDN, or dCDN is unable to pass a trigger request on to any of its affected dCDNs; dCDN should report an error if the request is abandoned, otherwise it must keep the trigger in state "pending" or "active" until it's acted upon or uCDN chooses to cancel it. Or, if the request is queued and dCDN will not report further status, dCDN may report the trigger as "complete" with an "etime" in the future.

Note that an "invalidate" trigger may be reported as "complete" when surrogates that may have the data are offline, if those surrogates will not use the affected data without first revalidating it when they are back online. This does not apply to "preposition" or "purge" triggers.

5. Properties of Triggers

5.1. Properties of Trigger Requests

Properties of Trigger Requests are defined in the following subsections.

Property: type

Description: This property defines the type of the trigger:

Type: TriggerType

Mandatory: Yes

Property: metadata.urls

Description: The uCDN URL for the metadata the trigger applies to.

Type: URLs

Mandatory: No, but at least one of 'metadata.*' or 'content.*' MUST be present and non-empty.

Property: content.urls

Description: URLs of content data the trigger applies to, see Section 5.1.1.

Type: URLs

Mandatory: No, but at least one of 'metadata.*' or 'content.*' MUST be present and non-empty.

Property: content.ccid

Description: The Content Collection IDentifier of data the trigger applies to.

Type: List of strings

Mandatory: No, but at least one of 'metadata.*' or 'content.*' MUST be present and non-empty.

Property: metadata.patterns

Description: The metadata the trigger applies to.

Type: List of PatternMatch

Mandatory: No, but at least one of 'metadata.*' or 'content.*' MUST be present and non-empty, and metadata.patterns MUST NOT be present if the TriggerType is Preposition.

Property: content.patterns

Description: The content data the trigger applies to.

Type: List of PatternMatch

Mandatory: No, but at least one of 'metadata.*' or 'content.*' MUST be present and non-empty, and content.patterns MUST NOT be present if the TriggerType is Preposition.

5.1.1. Content URLs

To refer to content in dCDN, uCDN must present URLs in the same form clients will use to access content in that dCDN, after transformation to remove any surrogate-specific parts of a 302-redirect URL form. By definition, it is always possible to locate content based on URLs in this form.

If content URLs are transformed by an intermediate CDN in a cascade, that intermediate CDN must transform URLs in trigger requests it passes to its dCDN.

When processing trigger requests, CDNs SHOULD ignore the URL scheme (http or https) in comparing URLs. For example, for an invalidate or purge trigger, content may be invalidated or purged regardless of the protocol clients use to request it.

5.2. Properties of Trigger Status Resources

Property: trigger

Description: The properties of trigger request that created this record.

Type: TriggerRequest

Mandatory: Yes

Property: ctime

Description: Time at which the request was received by dCDN.

Time is local to dCDN, there is no requirement to synchronise clocks between interconnected CDNs.

Type: AbsoluteTime

Mandatory: Yes

Property: mtime
Description: Time at which the resource was last modified.
Time is local to dCDN, there is no requirement to synchronise clocks between interconnected CDNs.
Type: AbsoluteTime
Mandatory: Yes

Property: etime
Description: Estimate of the time at which dCDN expects to complete the activity. Time is local to dCDN, there is no requirement to synchronise clocks between interconnected CDNs.
Type: AbsoluteTime
Mandatory: No

Property: status
Description: Current status of the triggered activity.
Type: TriggerStatus
Mandatory: Yes

Property: errors
Description: List of ErrorDesc.
Mandatory: No.

5.3. Properties of ErrorDesc

An ErrorDesc object is used to report failure for URLs and patterns in a trigger request.

Property: error
Type: ErrorCode.
Mandatory: Yes.
Description: List of metadata.urls, content.urls, metadata.patterns, content.patterns

Description: Metadata and content references copied from the trigger request. Only those URLs and patterns to which the error applies shall be included in each property, but those URLs and patterns shall be exactly as they appear in the request, dCDN must not generalise the URLs. (For example, if uCDN requests prepositioning of URLs "http://ucdn.example.com/a" and "http://ucdn.example.com/b", dCDN may not generalise its error report to Pattern "http://ucdn.example.com/*").
Mandatory: At least one of these properties is mandatory in each ErrorDesc.

Property: description
Description: A String containing a human-readable description of the error.

Mandatory: No.

5.4. Properties of Trigger Collections

Property: links

Description: References to Trigger Status Resources in the collection.

Type: List of Relationships.

Mandatory: Yes

Property: staleresourcetime

Description: The length of time for which dCDN guarantees to keep a completed Trigger Status Resource. After this time, dCDN MAY delete the resource and all references to it from collections.

Type: Integer, time in seconds.

Mandatory: Yes, in the collection of all Trigger Status Resources if dCDN deletes stale entries. If the property is present in the filtered collections, it MUST have the same value as in the collection of all Trigger Status Resources.

5.5. Trigger Resource Simple Data Type Descriptions

This section describes the simpler data types that are used for properties of Trigger Status resources.

5.5.1. TriggerType

This type defines the type of action being triggered, permitted actions are:

- o Preposition - a request for dCDN to acquire metadata or content.
- o Invalidate - a request for dCDN to invalidate metadata or content. After servicing this request the dCDN will not use the specified data without first re-validating it using, for example, an "If-None-Match" HTTP request. The dCDN need not erase the associated data.
- o Purge - a request for dCDN to erase metadata or content. After servicing the request, the specified data must not be held on dCDN.

5.5.2. TriggerStatus

This type describes the current status of a Trigger, possible values are:

- o Pending - the trigger has not yet been acted upon.
- o Active - the trigger is currently being acted upon.

- o Complete - the triggered activity completed successfully, or the trigger has been accepted and no further status update will be made.
- o Failed - the triggered activity could not be completed.

5.5.3. URLs

This type describes a set of references to metadata or content, it is simply a list of absolute URLs.

5.5.4. AbsoluteTime

Times are expressed in seconds since the UNIX epoch.

5.5.5. ErrorCode

This type is used by dCDN to report failures in trigger processing.

- o EMETA - dCDN was unable to acquire metadata required to fulfil the request.
- o ECONTENT - dCDN was unable to acquire content (preposition triggers only).
- o EPERM - uCDN does not have permission to trigger the requested activity (for example, the data is owned by another CDN).
- o EREJECT - dCDN is not willing to fulfil the request (for example, a preposition request for content at a time when dCDN would not accept Request Routing requests from uCDN).
- o ECDN - An internal error in dCDN or one of its downstream CDNs.

6. JSON Encoding of Objects

This encoding is based on that described in [I-D.ietf-cdni-metadata], but has been reproduced here while metadata work is in progress. Once that work is complete, the authors would look to align with the structure of the metadata draft and make reference to common definitions as appropriate.

The encoding for a CI/T object is a JSON object containing a dictionary of (key,value) pairs where the keys are the property names, and the values are the associated property values.

The keys of the dictionary are the names of the properties associated with the object and are therefore dependent on the specific object being encoded (i.e. dependent on the MIME Media Type of the returned resource). Likewise, the values associated with each key are dependent on the specific object being encoded (i.e. dependent on the MIME Media Type of the returned resource).

The "trigger" property of the top level JSON object lists the requested action.

Key: trigger

Description: An object specifying the trigger type and a set of data to act upon.

Type: A JSON object.

Mandatory: Yes.

Object keys in JSON are case sensitive and therefore any dictionary key defined by this document (for example the names of CI/T object properties) MUST always be represented in lowercase.

In addition to the properties of an object, the following additional keys may be present.

Key: base

Description: Provides a prefix for any relative URLs in the object. This is similar to the XML base tag [XML-BASE]. If absent, all URLs in the remainder of the document must be absolute URLs.

Type: URI

Mandatory: No

Key: links

Description: The relationships of this object to other addressable objects.

Type: Array of Relationships.

Mandatory: Yes

6.1. JSON Encoding of Embedded Types

6.1.1. TriggerType

Key: type

Description: One of "preposition", "invalidate" or "purge".

Type: string

6.1.2. TriggerStatus

Key: status

Description: One of "pending", "active", "failed", "complete"

Type: string

6.1.3. PatternMatch

A PatternMatch is encoded as a JSON Object containing a string to match and flags describing the type of match.

Key: pattern

Description: A pattern for string matching. The pattern may contain the wildcards * and ?, where * matches any sequence of characters (including the empty string) and ? matches exactly one character. The three literals \ , * and ? should be escaped as \\, * and \?

Type: String

Mandatory: Yes

Key: case-sensitive

Description: Flag indicating whether or not case-sensitive matching should be used.

Type: Boolean

Mandatory: No, default is case-insensitive match.

Key: match-query-string

Description: Flag indicating whether or not the query string should be included in the pattern match.

Type: Boolean

Mandatory: No, default is not to include query.

Example of case-sensitive prefix match against

```
"http://www.example.com/trailers/":  
{  
  "pattern": "http://www.example.com/trailers/*",  
  "case-sensitive": true  
}
```

6.1.4. ErrorDesc

ErrorDesc shall be encoded as a JSON object with the following keys:

Key: error

Type: ErrorCode

Mandatory: Yes

Keys: metadata.urls, content.urls

Type: Array of strings

Mandatory: At least one of metadata.* or content.* must be present.

Keys: metadata.patterns, content.patterns

Type: Array of PatternMatch

Mandatory: At least one of metadata.* or content.* must be present.

Key: description

Type: String

Mandatory: No.

6.1.5. ErrorCode

One of the strings "EMETA", "ECONTENT", "EPERM", "EREJECT" or "ECDN".

6.1.6. Relationship

JSON: A dictionary with the following keys:

- o href - The URI of the of the addressable object being referenced.
- o rel - The Relationship between the referring object and the object it is referencing.
- o type - The MIME Media Type of the referenced object. See Section 6.2 for the MIME Media Types of objects specified in this document.
- o title - An optional title for the Relationship/link.

Note: The above structure follows the pattern of atom:link in [RFC4287].

Example Relationship to a CI/T Resource within a CI/T Collection:

```
{
  "href": "http://triggers.cdni.example.com/trigger/12345",
  "rel": "Trigger",
  "type": "application/cdni.ci.TriggerStatus+json"
}
```

The format of relationship values is expected to align with other CDNI interfaces. For example, rather than use simple names (like "Trigger" in this case), there may be a namespace that allows well-known and proprietary values to co-exist.

6.2. MIME Media Types

Table 1 lists the MIME Media Type for the trigger request, and each trigger object (resource) that is retrievable through the CI/T interface.

Data Object	MIME Media Type
TriggerRequest	application/cdni.ci.TriggerRequest+json
TriggerStatus	application/cdni.ci.TriggerStatus+json
TriggerCollection	application/cdni.ci.TriggerCollection+json

Table 1: MIME Media Types for CDNI Trigger resources

7. Examples

The following sections provide examples of different CI/T objects encoded as JSON.

No authentication is shown in the following illustrative examples, it is anticipated that authentication mechanisms will be aligned with other CDNI Interfaces as and when those mechanisms are defined.

Discovery of the triggers interface is out of scope of this document. In an implementation, all URLs are under control of dCDN and the uCDN must not attempt to ascribe any meaning to individual elements of the path. In examples in this section, the following URLs are used as the location of the collections of triggers:

- o Collection of all Triggers belonging to one uCDN:
http://dcdn.example.com/triggers
- o Filtered collections:
 - Pending: http://dcdn.example.com/triggers/pending
 - Active: http://dcdn.example.com/triggers/active
 - Complete: http://dcdn.example.com/triggers/complete
 - Failed: http://dcdn.example.com/triggers/failed

7.1. Creating Triggers

Examples of uCDN triggering activity in dCDN:

7.1.1. Preposition

An example of a preposition request, a POST to the "AllTriggers" collection.

Note that "metadata.patterns" and "content.patterns" are not allowed in a preposition trigger request.

REQUEST:

```
POST /triggers HTTP/1.1
User-Agent: example-user-agent/0.1
Host: dcdn.example.com
Accept: */*
Content-Type: application/cdni.ci.TriggerRequest+json
Content-Length: 315

{
  "trigger" : {
    "type": "preposition",

    "metadata.urls" : [ "http://metadata.example.com/a/b/c" ],
    "content.urls" : [
      "http://www.example.com/a/b/c/1",
      "http://www.example.com/a/b/c/2",
      "http://www.example.com/a/b/c/3",
      "http://www.example.com/a/b/c/4"
    ]
  }
}
```

```
    ]  
  }  
}
```

RESPONSE:

```
HTTP/1.1 201 Created  
Date: Sat, 23 Feb 2013 14:20:06 GMT  
Content-Length: 472  
Content-Type: application/cdni.ci.TriggerStatus+json  
Location: http://dcdn.example.com/triggers/0  
Server: example-server/0.1
```

```
{  
  "ctime": 1361629206,  
  "etime": 1361629214,  
  "mtime": 1361629206,  
  "status": "pending",  
  "trigger": {  
    "content.urls": [  
      "http://www.example.com/a/b/c/1",  
      "http://www.example.com/a/b/c/2",  
      "http://www.example.com/a/b/c/3",  
      "http://www.example.com/a/b/c/4"  
    ],  
    "metadata.urls": [  
      "http://metadata.example.com/a/b/c"  
    ],  
    "type": "preposition"  
  }  
}
```

7.1.2. Invalidate

An example of an invalidate request, another POST to the "AllTriggers" collection. This instructs dCDN to re-validate the content at "http://www.example.com/a/index.html", as well as any metadata and content whose URLs are prefixed by "http://metadata.example.com/a/b/" and "http://www.example.com/a/b/" respectively, using case-insensitive matching.

REQUEST:

```
POST /triggers HTTP/1.1  
User-Agent: example-user-agent/0.1  
Host: dcdn.example.com  
Accept: */*  
Content-Type: application/cdni.ci.TriggerRequest+json  
Content-Length: 352
```

```
{
  "trigger" : {
    "type": "invalidate",

    "metadata.patterns" : [
      { "pattern" : "http://metadata.example.com/a/b/*" }
    ],

    "content.urls" : [ "http://www.example.com/a/index.html" ],
    "content.patterns" : [
      { "pattern" : "http://www.example.com/a/b/*",
        "case-sensitive" : true
      }
    ]
  }
}
```

RESPONSE:

```
HTTP/1.1 201 Created
Date: Sat, 23 Feb 2013 14:20:08 GMT
Content-Length: 551
Content-Type: application/cdni.ci.TriggerStatus+json
Location: http://dcdn.example.com/triggers/1
Server: example-server/0.1
```

```
{
  "ctime": 1361629208,
  "etime": 1361629216,
  "mtime": 1361629208,
  "status": "pending",
  "trigger": {
    "content.patterns": [
      {
        "case-sensitive": true,
        "pattern": "http://www.example.com/a/b/*"
      }
    ],
    "content.urls": [
      "http://www.example.com/a/index.html"
    ],
    "metadata.patterns": [
      {
        "pattern": "http://metadata.example.com/a/b/*"
      }
    ],
    "type": "invalidate"
  }
}
```

```
}
```

7.2. Examining Trigger Status

Once triggers have been created, uCDN can check their status as shown in these examples.

7.2.1. Collection of All Triggers

The uCDN can fetch the set of all the triggers it has created and which have not yet been deleted or removed as expired. After creation of the "preposition" and "invalidate" triggers shown above, this collection might look as follows:

REQUEST:

```
GET /triggers HTTP/1.1
User-Agent: example-user-agent/0.1
Host: dcdn.example.com
Accept: */*
```

RESPONSE:

```
HTTP/1.1 200 OK
Content-Length: 422
Expires: Sat, 23 Feb 2013 14:21:08 GMT
Server: example-server/0.1
ETag: "1484827667515030767"
Cache-Control: max-age=60
Date: Sat, 23 Feb 2013 14:20:08 GMT
Content-Type: application/cdni.ci.TriggerCollection+json

{
  "links": [
    {
      "href": "http://dcdn.example.com/triggers/0",
      "rel": "Trigger",
      "type": "application/cdni.ci.TriggerStatus+json"
    },
    {
      "href": "http://dcdn.example.com/triggers/1",
      "rel": "Trigger",
      "type": "application/cdni.ci.TriggerStatus+json"
    }
  ],
  "staleresourcetime": 86400
}
```

7.2.2. Filtered Collections of Triggers

The filtered collections are also available to uCDN. Before dCDN starts processing the two triggers shown above, both will appear in the collection of Pending Triggers, for example:

REQUEST:

```
GET /triggers/pending HTTP/1.1
User-Agent: example-user-agent/0.1
Host: dcdn.example.com
Accept: */*
```

RESPONSE:

```
HTTP/1.1 200 OK
Content-Length: 422
Expires: Sat, 23 Feb 2013 14:21:09 GMT
Server: example-server/0.1
ETag: "-970375801048973013"
Cache-Control: max-age=60
Date: Sat, 23 Feb 2013 14:20:09 GMT
Content-Type: application/cdni.ci.TriggerCollection+json
```

```
{
  "links": [
    {
      "href": "http://dcdn.example.com/triggers/0",
      "rel": "Trigger",
      "type": "application/cdni.ci.TriggerStatus+json"
    },
    {
      "href": "http://dcdn.example.com/triggers/1",
      "rel": "Trigger",
      "type": "application/cdni.ci.TriggerStatus+json"
    }
  ],
  "staleresourcetime": 86400
}
```

At this point, if no other triggers had been created, the other filtered views of the triggers would be empty. For example:

REQUEST:

```
GET /triggers/complete HTTP/1.1
User-Agent: example-user-agent/0.1
Host: dcdn.example.com
Accept: */*
```

RESPONSE:

```
HTTP/1.1 200 OK
Content-Length: 53
Expires: Sat, 23 Feb 2013 14:21:09 GMT
Server: example-server/0.1
ETag: "-654105208640281650"
Cache-Control: max-age=60
Date: Sat, 23 Feb 2013 14:20:09 GMT
Content-Type: application/cdni.ci.TriggerCollection+json
```

```
{
  "links": [],
  "staleresourcetime": 86400
}
```

7.2.3. Trigger Status Resources

The Trigger Status Resources can also be examined for detail about individual triggers. For example, for the "preposition" and "invalidate" triggers from previous examples:

REQUEST:

```
GET /triggers/0 HTTP/1.1
User-Agent: example-user-agent/0.1
Host: dcdn.example.com
Accept: */*
```

RESPONSE:

```
HTTP/1.1 200 OK
Content-Length: 472
Expires: Sat, 23 Feb 2013 14:21:08 GMT
Server: example-server/0.1
ETag: "-7651038857765989381"
Cache-Control: max-age=60
Date: Sat, 23 Feb 2013 14:20:08 GMT
Content-Type: application/cdni.ci.TriggerStatus+json
```

```
{
  "ctime": 1361629206,
  "etime": 1361629214,
  "mtime": 1361629206,
  "status": "pending",
  "trigger": {
    "content.urls": [
      "http://www.example.com/a/b/c/1",
      "http://www.example.com/a/b/c/2",
      "http://www.example.com/a/b/c/3",
      "http://www.example.com/a/b/c/4"
    ],
    "metadata.urls": [
      "http://metadata.example.com/a/b/c"
    ],
    "type": "preposition"
  }
}
```

REQUEST:

```
GET /triggers/1 HTTP/1.1
User-Agent: example-user-agent/0.1
Host: dcdn.example.com
Accept: */*
```

RESPONSE:

```
HTTP/1.1 200 OK
Content-Length: 551
Expires: Sat, 23 Feb 2013 14:21:09 GMT
Server: example-server/0.1
ETag: "-1103964172288811711"
Cache-Control: max-age=60
Date: Sat, 23 Feb 2013 14:20:09 GMT
Content-Type: application/cdni.ci.TriggerStatus+json
```

```
{
  "ctime": 1361629208,
  "etime": 1361629216,
  "mtime": 1361629208,
  "status": "pending",
  "trigger": {
    "content.patterns": [
      {
        "case-sensitive": true,
        "pattern": "http://www.example.com/a/b/*"
      }
    ],
    "content.urls": [
      "http://www.example.com/a/index.html"
    ],
    "metadata.patterns": [
      {
        "pattern": "http://metadata.example.com/a/b/*"
      }
    ],
    "type": "invalidate"
  }
}
```

7.2.4. Polling for Change

The uCDN may use the Entity Tags of collections or resources when polling for change in status, as shown in the following examples:

REQUEST:

```
GET /triggers/pending HTTP/1.1
User-Agent: example-user-agent/0.1
Host: dcdn.example.com
Accept: */*
If-None-Match: "-970375801048973013"
```

RESPONSE:

```
HTTP/1.1 304 Not Modified
Content-Length: 0
Expires: Sat, 23 Feb 2013 14:21:09 GMT
Server: example-server/0.1
ETag: "-970375801048973013"
Cache-Control: max-age=60
Date: Sat, 23 Feb 2013 14:20:09 GMT
Content-Type: application/cdni.ci.TriggerCollection+json
```

REQUEST:

```
GET /triggers/0 HTTP/1.1
User-Agent: example-user-agent/0.1
Host: dcdn.example.com
Accept: */*
If-None-Match: "-7651038857765989381"
```

RESPONSE:

```
HTTP/1.1 304 Not Modified
Content-Length: 0
Expires: Sat, 23 Feb 2013 14:21:08 GMT
Server: example-server/0.1
ETag: "-7651038857765989381"
Cache-Control: max-age=60
Date: Sat, 23 Feb 2013 14:20:08 GMT
Content-Type: application/cdni.ci.TriggerStatus+json
```

When the triggered activity is complete, the contents of the filtered collections will be updated, along with their Entity Tags. For example, when the two example triggers are complete, the collections of pending and complete triggers may look like:

REQUEST:

```
GET /triggers/pending HTTP/1.1
User-Agent: example-user-agent/0.1
Host: dcdn.example.com
Accept: */*
If-None-Match: "-970375801048973013"
```

RESPONSE:

```
HTTP/1.1 200 OK
Content-Length: 53
Expires: Sat, 23 Feb 2013 14:21:13 GMT
Server: example-server/0.1
ETag: "-7056231826368088123"
Cache-Control: max-age=60
Date: Sat, 23 Feb 2013 14:20:13 GMT
Content-Type: application/cdni.ci.TriggerCollection+json

{
  "links": [],
  "staleresourcetime": 86400
}
```

REQUEST:

```
GET /triggers/complete HTTP/1.1
User-Agent: example-user-agent/0.1
Host: dcdn.example.com
Accept: */*
```

RESPONSE:

```
HTTP/1.1 200 OK
Content-Length: 422
Expires: Sat, 23 Feb 2013 14:21:20 GMT
Server: example-server/0.1
ETag: "2013095476705515794"
Cache-Control: max-age=60
Date: Sat, 23 Feb 2013 14:20:20 GMT
Content-Type: application/cdni.ci.TriggerCollection+json
```

```
{
  "links": [
    {
      "href": "http://dcdn.example.com/triggers/0",
      "rel": "Trigger",
      "type": "application/cdni.ci.TriggerStatus+json"
    },
    {
      "href": "http://dcdn.example.com/triggers/1",
      "rel": "Trigger",
      "type": "application/cdni.ci.TriggerStatus+json"
    }
  ],
  "staleresourcetime": 86400
}
```

7.2.5. Cancelling or Removing a Trigger

To request dCDN to cancel a Trigger, uCDN may delete the Trigger Resource. It may also delete completed and failed triggers to reduce the size of the collections. For example, to remove the "preposition" request from earlier examples:

REQUEST:

```
DELETE /triggers/0 HTTP/1.1
User-Agent: example-user-agent/0.1
Host: dcdn.example.com
Accept: */*
```

RESPONSE:

```
HTTP/1.1 204 No Content
Date: Sat, 23 Feb 2013 14:20:20 GMT
Content-Length: 0
Content-Type: text/html; charset=UTF-8
Server: example-server/0.1
```

This would, for example, cause the collection of completed triggers shown in the example above to be updated to:

REQUEST:

```
GET /triggers/complete HTTP/1.1
User-Agent: example-user-agent/0.1
Host: dcdn.example.com
Accept: */*
```

RESPONSE:

```
HTTP/1.1 200 OK
Content-Length: 239
Expires: Sat, 23 Feb 2013 14:21:20 GMT
Server: example-server/0.1
ETag: "4257416552489354137"
Cache-Control: max-age=60
Date: Sat, 23 Feb 2013 14:20:20 GMT
Content-Type: application/cdni.ci.TriggerCollection+json
```

```
{
  "links": [
    {
      "href": "http://dcdn.example.com/triggers/1",
      "rel": "Trigger",
      "type": "application/cdni.ci.TriggerStatus+json"
    }
  ],
  "staleresourcetime": 86400
}
```

7.2.6. Error Reporting

In this example uCDN has requested prepositioning of "http://newsite.example.com/index.html", but dCDN was unable to locate metadata for that site:

REQUEST:

```
GET /triggers/2 HTTP/1.1
User-Agent: example-user-agent/0.1
Host: dcdn.example.com
Accept: */*
```

RESPONSE:

```
HTTP/1.1 200 OK
Content-Length: 505
Expires: Sat, 23 Feb 2013 14:21:28 GMT
Server: example-server/0.1
ETag: "2621489144226897896"
Cache-Control: max-age=60
Date: Sat, 23 Feb 2013 14:20:28 GMT
Content-Type: application/cdni.ci.TriggerStatus+json
```

```
{
  "ctime": 1361629220,
  "errors": [
    {
      "content.urls": [
        "http://newsite.example.com/index.html"
      ],
      "description":
        "No HostIndex entry found for newsite.example.com",
      "error": "EMETA"
    }
  ],
  "etime": 1361629228,
  "mtime": 1361629224,
  "status": "active",
  "trigger": {
    "content.urls": [
      "http://newsite.example.com/index.html"
    ],
    "type": "preposition"
  }
}
```

8. IANA Considerations

TBD.

9. Security Considerations

The dCDN must ensure that each uCDN only has access to its own Trigger Status Resources.

It is anticipated that a common authentication mechanism will be used by this and other CDNI Interconnect interfaces, the mechanism must exist but is not identified in this document.

The dCDN must ensure that activity triggered by uCDN only affects metadata or content originating from that uCDN.

10. Acknowledgements

TBD.

11. References

11.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2616, June 1999.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005.

11.2. Informative References

- [I-D.ietf-cdni-framework] Peterson, L. and B. Davie, "Framework for CDN Interconnection", draft-ietf-cdni-framework-05 (work in progress), September 2013.
- [I-D.ietf-cdni-metadata] Niven-Jenkins, B., Murray, R., Watson, G., Caulfield, M., Leung, K., and K. Ma, "CDN Interconnect Metadata",

draft-ietf-cdni-metadata-02 (work in progress), July 2013.

[I-D.ietf-cdni-requirements]

Leung, K. and Y. Lee, "Content Distribution Network Interconnection (CDNI) Requirements", draft-ietf-cdni-requirements-10 (work in progress), September 2013.

[RFC4287] Nottingham, M., Ed. and R. Sayre, Ed., "The Atom Syndication Format", RFC 4287, December 2005.

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

[XML-BASE]

Marsh, J., Ed. and R. Tobin, Ed., "XML Base (Second Edition) - <http://www.w3.org/TR/xmlbase/>", January 2009.

Authors' Addresses

Rob Murray
Velocix (Alcatel-Lucent)
3 Ely Road
Milton, Cambridge CB24 6DD
UK

Email: rmurray@velocix.com

Ben Niven-Jenkins
Velocix (Alcatel-Lucent)
3 Ely Road
Milton, Cambridge CB24 6DD
UK

Email: ben@velocix.com

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: November 20, 2016

R. Murray
B. Niven-Jenkins
Nokia
May 19, 2016

CDNI Control Interface / Triggers
draft-ietf-cdni-control-triggers-15

Abstract

This document describes the part of the CDN Interconnection Control Interface that allows a CDN to trigger activity in an interconnected CDN that is configured to deliver content on its behalf. The upstream CDN can use this mechanism to request that the downstream CDN pre-positions metadata or content, or that it invalidates or purges metadata or content. The upstream CDN can monitor the status of activity that it has triggered in the downstream CDN.

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 November 20, 2016.

Copyright Notice

Copyright (c) 2016 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. Model for CDNI Triggers	4
2.1. Timing of Triggered Activity	6
2.2. Scope of Triggered Activity	6
2.2.1. Multiple Interconnected CDNs	7
2.3. Trigger Results	8
3. Collections of Trigger Status Resources	8
4. CDNI Trigger Interface	9
4.1. Creating Triggers	10
4.2. Checking Status	11
4.2.1. Polling Trigger Status Resource collections	12
4.2.2. Polling Trigger Status Resources	12
4.3. Cancelling Triggers	12
4.4. Deleting Triggers	13
4.5. Expiry of Trigger Status Resources	14
4.6. Loop Detection and Prevention	14
4.7. Error Handling	15
4.8. Content URLs	16
5. CI/T Object Properties and Encoding	16
5.1. CI/T Objects	16
5.1.1. CI/T Commands	16
5.1.2. Trigger Status Resource	17
5.1.3. Trigger Collection	19
5.2. Properties of CI/T Objects	20
5.2.1. Trigger Specification	20
5.2.2. Trigger Type	21
5.2.3. Trigger Status	22
5.2.4. PatternMatch	23
5.2.5. Absolute Time	24
5.2.6. Error Description	24
5.2.7. Error Code	25
6. Examples	26
6.1. Creating Triggers	26
6.1.1. Preposition	26
6.1.2. Invalidate	28

6.2.	Examining Trigger Status	29
6.2.1.	Collection of All Triggers	29
6.2.2.	Filtered Collections of Trigger Status Resources	30
6.2.3.	Individual Trigger Status Resources	32
6.2.4.	Polling for Change	34
6.2.5.	Deleting Trigger Status Resources	37
6.2.6.	Error Reporting	38
7.	IANA Considerations	39
7.1.	CDNI Payload Type Parameter Registrations	40
7.2.	CDNI CI/T Trigger Types Registry	40
7.3.	CDNI CI/T Error Codes Registry	40
8.	Security Considerations	41
8.1.	Authentication, Authorization, Confidentiality, Integrity Protection	41
8.2.	Denial of Service	42
8.3.	Privacy	43
9.	Acknowledgements	43
10.	References	43
10.1.	Normative References	43
10.2.	Informative References	44
Appendix A.	Formalization of the JSON Data	45
Authors' Addresses	46

1. Introduction

[RFC6707] introduces the problem scope for CDN Interconnection (CDNI) and lists the four categories of interfaces that may be used to compose a CDNI solution (Control, Metadata, Request Routing, Logging).

[RFC7336] expands on the information provided in [RFC6707] and describes each of the interfaces and the relationships between them in more detail.

This document describes the "CI/T" interface, "CDNI Control interface / Triggers". It does not consider those parts of the control interface that relate to configuration, bootstrapping or authentication of CDN Interconnect interfaces. Section 4 of [RFC7337] identifies the requirements specific to the CI interface, requirements applicable to the CI/T interface are CI-1 to CI-6.

- o Section 2 outlines the model for the CI/T Interface at a high level.
- o Section 3 describes collections of Trigger Status Resources.
- o Section 4 defines the web service provided by the dCDN.

- o Section 5 lists properties of CI/T Commands and Status Resources.
- o Section 6 contains example messages.

1.1. Terminology

This document reuses the terminology defined in [RFC6707] and uses "uCDN" and "dCDN" as shorthand for "Upstream CDN" and "Downstream CDN", respectively.

2. Model for CDNI Triggers

A CI/T Command, sent from the uCDN to the dCDN, is a request for the dCDN to do some work relating to data associated with content requests originating from the uCDN.

There are two types of CI/T Command: CI/T Trigger Commands, and CI/T Cancel Commands. The CI/T Cancel Command can be used to request cancellation of an earlier CI/T Trigger Command. A CI/T Trigger Command is of one of the following types:

- o preposition - used to instruct the dCDN to fetch metadata from the uCDN, or content from any origin including the uCDN.
- o invalidate - used to instruct the dCDN to revalidate specific metadata or content before re-using it.
- o purge - used to instruct the dCDN to delete specific metadata or content.

The CI/T interface is a web service offered by the dCDN. It allows CI/T commands to be issued, and triggered activity to be tracked. The CI/T interface builds on top of HTTP/1.1 [RFC7230]. References to URL in this document relate to http/https URIs, as defined in [RFC7230] section 2.7.

When the dCDN accepts a CI/T Command it creates a resource describing status of the triggered activity, a Trigger Status Resource. The uCDN can poll Trigger Status Resources to monitor progress.

The dCDN maintains at least one collection of Trigger Status Resources for each uCDN. Each uCDN only has access to its own collections, the locations of which are shared when CDN interconnection is established.

To trigger activity in the dCDN, the uCDN POSTs a CI/T Command to the collection of Trigger Status Resources. If the dCDN accepts the CI/T Command, it creates a new Trigger Status Resource and returns its

location to the uCDN. To monitor progress, the uCDN can GET the Trigger Status Resource. To request cancellation of a CI/T Trigger Command the uCDN can POST to the collection of Trigger Status Resources, or simply DELETE the Trigger Status Resource.

In addition to the collection of all Trigger Status Resources for the uCDN, the dCDN can maintain filtered views of that collection. These filtered views are defined in Section 3 and include collections of Trigger Status Resources corresponding to active and completed CI/T Trigger Commands. These collections provide a mechanism for polling the status of multiple jobs.

Figure 1 is an example showing the basic message flow used by the uCDN to trigger activity in the dCDN, and for the uCDN to discover the status of that activity. Only successful triggering is shown. Examples of the messages are given in Section 6.

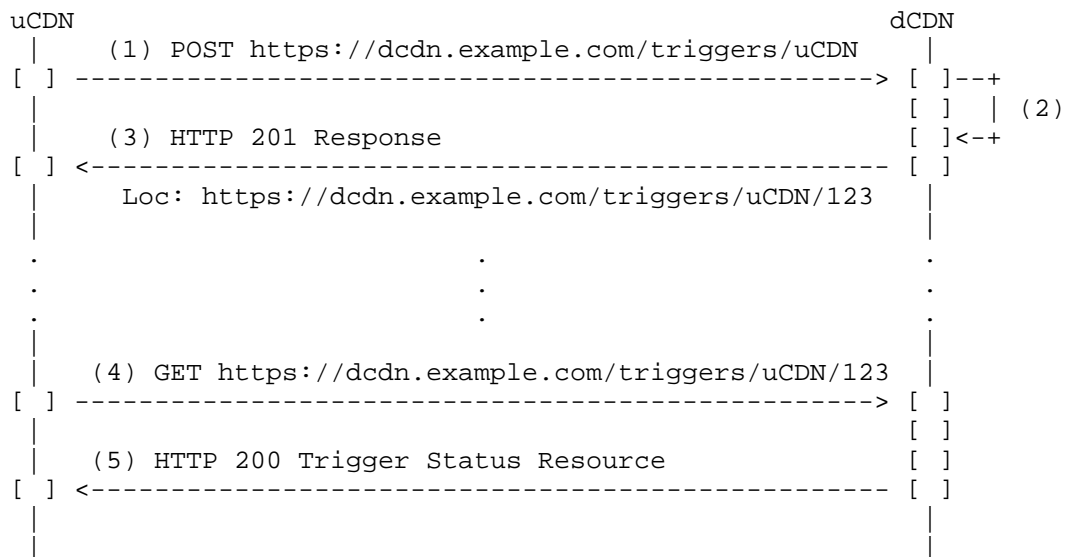


Figure 1: Basic CDNI Message Flow for Triggers

The steps in Figure 1 are:

1. The uCDN triggers action in the dCDN by posting a CI/T Command to a collection of Trigger Status Resources, "https://dcdn.example.com/triggers/uCDN". The URL of this was given to the uCDN when the CI/T interface was established.

2. The dCDN authenticates the request, validates the CI/T Command and, if it accepts the request, creates a new Trigger Status Resource.
3. The dCDN responds to the uCDN with an HTTP 201 response status, and the location of the Trigger Status Resource.
4. The uCDN can poll, possibly repeatedly, the Trigger Status Resource in the dCDN.
5. The dCDN responds with the Trigger Status Resource, describing progress or results of the CI/T Trigger Command.

The remainder of this document describes the messages, Trigger Status Resources, and collections of Trigger Status Resources in more detail.

2.1. Timing of Triggered Activity

Timing of the execution of CI/T Commands is under the dCDN's control, including its start-time and pacing of the activity in the network.

CI/T invalidate and purge commands MUST be applied to all data acquired before the command was accepted by the dCDN. The dCDN SHOULD NOT apply CI/T invalidate and purge commands to data acquired after the CI/T Command was accepted, but this may not always be achievable so the uCDN cannot count on that.

If the uCDN wishes to invalidate or purge content then immediately pre-position replacement content at the same URLs, it SHOULD ensure the dCDN has completed the invalidate/purge before initiating the prepositioning. Otherwise, there is a risk that the dCDN pre-positions the new content, then immediately invalidates or purges it (as a result of the two uCDN requests running in parallel).

Because the CI/T Command timing is under the dCDN's control, the dCDN implementation can choose whether to apply CI/T invalidate and purge commands to content acquisition that has already started when the command is received.

2.2. Scope of Triggered Activity

Each CI/T Command can operate on multiple metadata and content URLs.

Multiple representations of an HTTP resource may share the same URL. CI/T Trigger Commands that invalidate or purge metadata or content apply to all resource representations with matching URLs.

2.2.1. Multiple Interconnected CDNs

In a network of interconnected CDNs a single uCDN will originate a given item of metadata and associated content, it may distribute that metadata and content to more than one dCDN, which may in-turn distribute that metadata and content to further-downstream CDNs.

An intermediate CDN is a dCDN that passes on CDNI metadata and content to further-downstream dCDNs.

A diamond configuration is one where a dCDN can acquire metadata and content originated in one uCDN from that uCDN itself and an intermediate CDN, or via more than one intermediate CDN.

CI/T commands originating in the single source uCDN affect metadata and content in all dCDNs but, in a diamond configuration, it may not be possible for the dCDN to determine which uCDN it acquired content from. In this case a dCDN MUST allow each uCDN from which it may have acquired the content to act upon that content using CI/T Commands.

In all other cases, a dCDN MUST reject CI/T Commands from a uCDN that act on another uCDN's data using, for example, HTTP "403 Forbidden".

Security considerations are discussed further in Section 8.

The diamond configuration may lead to inefficient interactions, but the interactions are otherwise harmless. For example:

- o When the uCDN issues an invalidate CI/T command, a dCDN will receive that command from multiple directly connected uCDNs. The dCDN may schedule multiple those commands separately, and the last may affect content already revalidated following execution of the invalidate command scheduled first.
- o If one of a dCDN's directly-connected uCDNs loses its rights to distribute content, it may issue a CI/T purge command. That purge may affect content the dCDN could retain because it's distributed by another directly-connected uCDN. But, that content can be re-acquired by the dCDN from the remaining uCDN.
- o When the uCDN originating an item of content issues a CI/T purge followed by a preposition - two directly connected uCDNs will pass those commands to a dCDN. That dCDN implementation need not merge those operations, or notice the repetition. In which case the purge issued by one uCDN will complete before the other. The first uCDN to finish its purge may then forward the preposition trigger, and content pre-positioned as a result might be affected

by the still-running purge issued by the other uCDN. However, the dCDN will re-acquire that content as needed, or when it's asked to pre-position the content by the second uCDN. A dCDN implementation could avoid this interaction by knowing which uCDN it acquired the content from, or it could minimize the consequences by recording the time at which the invalidate/purge command was received and not applying it to content acquired after that time.

2.3. Trigger Results

Possible states for a Trigger Status Resource are defined in section Section 5.2.3.

The CI/T Trigger Command MUST NOT be reported as 'complete' until all actions have been completed successfully. The reasons for failure, and URLs or Patterns affected, SHOULD be enumerated in the Trigger Status Resource. For more detail, see section Section 4.7.

If a dCDN is also acting as a uCDN in a cascade, it MUST forward CI/T Commands to any downstream CDNs that may be affected. The CI/T Trigger Command MUST NOT be reported as 'complete' in a CDN until it is 'complete' in all of its downstream CDNs. If a CI/T Trigger Command is reported as 'processed' in any dCDN, intermediate CDNs MUST NOT report 'complete', instead they MUST also report 'processed'. A CI/T Command MAY be reported as 'failed' as soon as it fails in a CDN or in any of its downstream CDNs. A cancelled CI/T Trigger Command MUST be reported as 'cancelling' until it has been reported as 'cancelled', 'complete', or 'failed' by all dCDNs in a cascade.

3. Collections of Trigger Status Resources

As described in Section 2, Trigger Status Resources exist in the dCDN to report the status of activity triggered by each uCDN.

A collection of Trigger Status Resources is a resource that contains a reference to each Trigger Status Resource in that collection.

The dCDN MUST make a collection of a uCDN's Trigger Status Resources available to that uCDN. This collection includes all of the Trigger Status Resources created for CI/T Commands from the uCDN that have been accepted by the dCDN, and have not yet been deleted by the uCDN, or expired and removed by the dCDN (as described in section Section 4.4). Trigger Status Resources belonging to a uCDN MUST NOT be visible to any other CDN. The dCDN could, for example, achieve this by offering different collection URLs to each uCDN, and by

filtering the response based on the uCDN with which the HTTP client is associated.

To trigger activity in a dCDN, or to cancel triggered activity, the uCDN POSTs a CI/T Command to the dCDN's collection of the uCDN's Trigger Status Resources.

In order to allow the uCDN to check the status of multiple jobs in a single request, the dCDN MAY also maintain collections representing filtered views of the collection of all Trigger Status Resources. These filtered collections are optional-to-implement but, if implemented, the dCDN MUST include links to them in the collection of all Trigger Status Resources. The filtered collections are:

- o Pending - Trigger Status Resources for CI/T Trigger Commands that have been accepted, but not yet acted upon.
- o Active - Trigger Status Resources for CI/T Trigger Commands that are currently being processed in the dCDN.
- o Complete - Trigger Status Resources representing activity that completed successfully, and 'processed' CI/T Trigger Commands for which no further status updates will be made by the dCDN.
- o Failed - Trigger Status Resources representing CI/T Commands that failed or were cancelled by the uCDN.

4. CDNI Trigger Interface

This section describes an interface to enable an upstream CDN to trigger activity in a downstream CDN.

The CI/T interface builds on top of HTTP, so dCDNs may make use of any HTTP feature when implementing the CI/T interface. For example, a dCDN SHOULD make use of HTTP's caching mechanisms to indicate that a requested response/representation has not been modified, reducing the uCDN's processing needed to determine whether the status of triggered activity has changed.

All dCDNs implementing CI/T MUST support the HTTP GET, HEAD, POST and DELETE methods as defined in [RFC7231].

The only representation specified in this document is JSON, [RFC7159]. It MUST be supported by the uCDN and by the dCDN.

The URL of the dCDN's collection of all Trigger Status Resources needs to be either discovered by, or configured in, the uCDN. The

mechanism for discovery of that URL is outside the scope of this document.

CI/T Commands are POSTed to the dCDN's collection of all Trigger Status Resources. If a CI/T Trigger Command is accepted by the dCDN, the dCDN creates a new Trigger Status Resource and returns its URI to the uCDN in an HTTP 201 response. The triggered activity can then be monitored by the uCDN using that resource and the collections described in Section 3.

The URI of each Trigger Status Resource is returned to the uCDN when it is created, and URIs of all Trigger Status Resources are listed in the dCDN's collection of all Trigger Status Resources. This means all Trigger Status Resources can be discovered by the uCDN, so dCDNs are free to assign whatever structure they desire to the URIs for CI/T resources. Therefore uCDNs MUST NOT make any assumptions regarding the structure of CI/T URIs or the mapping between CI/T objects and their associated URIs. URIs present in the examples in this document are purely illustrative and are not intended to impose a definitive structure on CI/T interface implementations.

4.1. Creating Triggers

To issue a CI/T Command, the uCDN makes an HTTP POST to the dCDN's collection of all of the uCDN's Trigger Status Resources. The request body of that POST is a CI/T Command, as described in Section 5.1.1.

The dCDN validates the CI/T Command. If the command is malformed or the uCDN does not have sufficient access rights, the dCDN MUST either respond with an appropriate 4xx HTTP error code and not create a Trigger Status Resource, or create a 'failed' Trigger Status Resource containing an appropriate error description.

When a CI/T Trigger Command is accepted, the uCDN MUST create a new Trigger Status Resource which will convey a specification of the CI/T Command and its current status. The HTTP response to the dCDN MUST have status code 201 and MUST convey the URI of the Trigger Status Resource in the Location header field. The HTTP response SHOULD include the content of the newly created Trigger Status Resource. This is particularly important in cases where the CI/T Trigger Command has completed immediately.

Once a Trigger Status Resource has been created the dCDN MUST NOT re-use its URI, even after that Trigger Status Resource has been removed.

The dCDN SHOULD track and report on progress of CI/T Trigger Commands using a Trigger Status Resource, Section 5.1.2. If the dCDN is not able to do that, it MUST indicate that it has accepted the request but will not be providing further status updates. To do this, it sets the status of the Trigger Status Resource to "processed". In this case, CI/T processing should continue as for a "complete" request, so the Trigger Status Resource MUST be added to the dCDN's collection of Complete Trigger Status Resources. The dCDN SHOULD also provide an estimated completion time for the request, by using the "etime" property of the Trigger Status Resource. This will allow the uCDN to schedule prepositioning after an earlier delete of the same URLs is expected to have finished.

If the dCDN is able to track the execution of CI/T Commands and a CI/T Command is queued by the dCDN for later action, the status property of the Trigger Status Resource MUST be "pending". Once processing has started the "status" MUST be "active". Finally, once the CI/T Command is complete, the status MUST be set to "complete" or "failed".

A CI/T Trigger Command may result in no activity in the dCDN if, for example, it is an invalidate or purge request for data the dCDN has not yet acquired, or a pre-position request for data it has already acquired and which is still valid. In this case, the "status" of the Trigger Status Resource MUST be "processed" or "complete", and the Trigger Status Resource MUST be added to the dCDN's collection of Complete Trigger Status Resources.

Once created, Trigger Status Resources can be cancelled or deleted by the uCDN, but not modified. The dCDN MUST reject PUT and POST requests from the uCDN to Trigger Status Resources by responding with an appropriate HTTP status code, for example 405 "Method Not Allowed".

4.2. Checking Status

The uCDN has two ways to check progress of CI/T Commands it has issued to the dCDN, described in sections Section 4.2.1 and Section 4.2.2.

To allow the uCDN to check for change in status of a Trigger Status Resource or collection of Trigger Status Resources without re-fetching the whole Resource or Collection, the dCDN SHOULD include Entity Tags for the uCDN to use as cache validators, as defined in [RFC7232].

The dCDN SHOULD use the cache control headers for responses to GETs for Trigger Status Resources and Collections to indicate the frequency at which it recommends the uCDN should poll for change.

4.2.1. Polling Trigger Status Resource collections

The uCDN can fetch the collection of its Trigger Status Resources, or filtered views of that collection.

This makes it possible to poll status of all CI/T Trigger Commands in a single request. If the dCDN moves a Trigger Status Resource from the Active to the Completed collection, the uCDN can fetch the result of that activity.

When polling in this way, the uCDN SHOULD use HTTP Entity Tags to monitor for change, rather than repeatedly fetching the whole collection. An example of this is given in section Section 6.2.4.

4.2.2. Polling Trigger Status Resources

The uCDN has a URI provided by the dCDN for each Trigger Status Resource it has created, it may fetch that Trigger Status Resource at any time.

This can be used to retrieve progress information, and to fetch the result of the CI/T Command.

When polling in this way, the uCDN SHOULD use HTTP Entity Tags to monitor for change, rather than repeatedly fetching the Trigger Status Resource.

4.3. Cancelling Triggers

The uCDN can request cancellation of a CI/T Trigger Command by POSTing a CI/T Cancel Command to the collection of all Trigger Status Resources.

The dCDN is required to accept and respond to the CI/T Cancel Command, but the actual cancellation of a CI/T Trigger Command is optional-to-implement.

The dCDN MUST respond to the CI/T Cancel Command appropriately, for example with HTTP status code 200 "OK" if the cancellation has been processed and the CI/T Command is inactive, 202 "Accepted" if the command has been accepted but the CI/T Command remains active, or 501 "Not Implemented" if cancellation is not supported by the dCDN.

If cancellation of a "pending" Trigger Status Resource is accepted by the dCDN, the dCDN SHOULD NOT start processing of that activity. Issuing a CT/T Cancel Command for a "pending" Trigger Status Resource does not however guarantee that the corresponding activity will not be started, because the uCDN cannot control the timing of that activity. Processing could, for example, start after the POST is sent by the uCDN but before that request is processed by the dCDN.

If cancellation of an "active" or "processed" Trigger Status Resource is accepted by the dCDN, the dCDN SHOULD stop processing the CI/T Command. However, as with cancellation of a "pending" CI/T Command, the dCDN does not guarantee this.

If the CI/T Command cannot be stopped immediately, the status in the corresponding Trigger Status Resource MUST be set to "cancelling", and the Trigger Status Resource MUST remain in the collection of Trigger Status Resources for active CI/T Commands. If processing is stopped before normal completion, the status value in the Trigger Status Resource MUST be set to "cancelled", and the Trigger Status Resource MUST be included in the collection of failed CT/T Trigger Commands.

Cancellation of a "complete" or "failed" Trigger Status Resource requires no processing in the dCDN. Its status MUST NOT be changed to "cancelled".

4.4. Deleting Triggers

The uCDN can delete Trigger Status Resources at any time, using the HTTP DELETE method. The effect is similar to cancellation, but no Trigger Status Resource remains afterwards.

Once deleted, the references to a Trigger Status Resource MUST be removed from all Trigger Status Resource collections. Subsequent requests to GET the deleted Trigger Status Resource SHOULD be rejected by the dCDN with an HTTP error.

If a "pending" Trigger Status Resource is deleted, the dCDN SHOULD NOT start processing of that activity. Deleting a "pending" Trigger Status Resource does not however guarantee that it has not started because the uCDN cannot control the timing of that activity. Processing may, for example, start after the DELETE is sent by the uCDN but before that request is processed by the dCDN.

If an "active" or "processed" Trigger Status Resource is deleted, the dCDN SHOULD stop processing the CI/T Command. However, as with deletion of a "pending" Trigger Status Resource, the dCDN does not guarantee this.

Deletion of a "complete" or "failed" Trigger Status Resource requires no processing in the dCDN other than deletion of the Trigger Status Resource.

4.5. Expiry of Trigger Status Resources

The dCDN can choose to automatically delete Trigger Status Resources some time after they become "complete", "processed", "failed" or "cancelled". In this case, the dCDN will remove the Trigger Status Resource and respond to subsequent requests for it with an HTTP error.

If the dCDN does remove Trigger Status Resources automatically, it MUST report the length of time after which it will do so, using a property of the collection of all Trigger Status Resources. It is RECOMMENDED that Trigger Status Resources are not automatically deleted by the dCDN for at least 24 hours after they become "complete", "processed", "failed" or "cancelled".

To ensure it is able to get the status of its Trigger Status Resources for completed and failed CI/T Commands, it is RECOMMENDED that the uCDN polling interval is less than the time after which records for completed activity will be deleted.

4.6. Loop Detection and Prevention

Given three CDNs, A, B and C, if CDNs B and C delegate delivery of CDN A's content to each other, CDN A's CI/T Commands could be passed between CDNs B and C in a loop. More complex networks of CDNs could contain similar loops involving more hops.

In order to prevent and detect such CI/T loops, each CDN uses a CDN Provider ID to uniquely identify itself. In every CI/T Command it originates or cascades, each CDN MUST append an array element containing its CDN Provider ID to a JSON array under an entry named "cdn-path". When receiving CI/T Commands a dCDN MUST check the cdn-path and reject any CI/T Command which already contains its own CDN Provider ID in the cdn-path. Transit CDNs MUST check the cdn-path and not cascade the CI/T Command to dCDNs that are already listed in cdn-path.

The CDN Provider Id consists of the two characters "AS" followed by the CDN Provider's Autonomous System number [RFC1930], then a colon (':') and an additional qualifier that is used to guarantee uniqueness in case a particular AS has multiple independent CDNs deployed. For example "AS64496:0".

If the CDN provider has multiple Autonomous Systems, the same AS number SHOULD be used in all messages from that CDN provider, unless there are multiple distinct CDNs.

If the RI interface described in [I-D.ietf-cdni-redirection] is implemented by the dCDN, the CI/T and RI interfaces SHOULD use the same CDN Provider Id.

4.7. Error Handling

A dCDN can signal rejection of a CI/T Command using HTTP status codes. For example, 400 if the request is malformed, or 403 or 404 if the uCDN does not have permission to issue CI/T Commands or it is trying to act on another CDN's data.

If any part of the CI/T Trigger Command fails, the trigger SHOULD be reported as "failed" once its activity is complete or if no further errors will be reported. The "errors" property in the Trigger Status Resource will be used to enumerate which actions failed and the reasons for failure, and can be present while the Trigger Status Resource is still "pending" or "active", if the CI/T Trigger Command is still running for some URLs or Patterns in the Trigger Specification.

Once a request has been accepted, processing errors are reported in the Trigger Status Resource using a list of Error Descriptions. Each Error Description is used to report errors against one or more of the URLs or Patterns in the Trigger Specification.

If a surrogate affected by a CI/T Trigger Command is offline in the dCDN, or the dCDN is unable to pass a CI/T Command on to any of its cascaded dCDNs:

- o If the CI/T Command is abandoned by the dCDN, the dCDN SHOULD report an error.
- o A CI/T "invalidate" command may be reported as "complete" when surrogates that may have the data are offline. In this case, surrogates MUST NOT use the affected data without first revalidating it when they are back online.
- o CI/T "preposition" and "purge" commands can be reported as "processed" if affected caches are offline and the activity will complete when they return to service.
- o Otherwise, the dCDN SHOULD keep the Trigger Status Resource in state "pending" or "active" until the CI/T Command is acted upon, or the uCDN chooses to cancel it.

4.8. Content URLs

If content URLs are transformed by an intermediate CDN in a cascade, that intermediate CDN MUST similarly transform URLs in CI/T Commands it passes to its dCDN.

When processing Trigger Specifications, CDNs MUST ignore the URL scheme (http or https) in comparing URLs. For example, for a CI/T invalidate or purge command, content MUST be invalidated or purged regardless of the protocol clients use to request it.

5. CI/T Object Properties and Encoding

The CI/T Commands, Trigger Status Resources and Trigger Collections and their properties are encoded using JSON, as defined in sections Section 5.1.1, Section 5.2.1, and Section 5.1.2. They MUST use the MIME Media Type 'application/cdni', with parameter 'ptype' values as defined below and in Section 7.1.

Names in JSON are case sensitive. The names and literal values specified in the present document MUST always use lower-case.

JSON types, including 'object', 'array', 'number' and 'string' are defined in [RFC7159].

Unrecognised name/value pairs in JSON objects SHOULD NOT be treated as an error by either the uCDN or dCDN. They SHOULD be ignored in the processing, and passed on by dCDN to any further dCDNs in a cascade.

5.1. CI/T Objects

The top-level objects defined by the CI/T interface are described in this section.

The encoding of values used by these objects is described in Section 5.2.

5.1.1. CI/T Commands

CI/T Commands MUST use a MIME Media Type of 'application/cdni; ptype=ci-trigger-command'.

A CI/T Command is encoded as a JSON object containing the following name/value pairs.

Name: trigger

Description: A specification of the trigger type, and a set of data to act upon.

Value: A Trigger Specification, as defined in Section 5.2.1.

Mandatory: No, but exactly one of "trigger" or "cancel" MUST be present in a CI/T Command.

Name: cancel

Description: The URLs of Trigger Status Resources for CI/T Trigger Commands that the uCDN wants to cancel.

Value: A non-empty JSON array of URLs represented as JSON strings.

Mandatory: No, but exactly one of "trigger" or "cancel" MUST be present in a CI/T Command.

Name: cdn-path

Description: The CDN Provider Identifiers of CDNs that have already issued the CI/T Command to their dCDNs.

Value: A non-empty JSON array of JSON strings, where each string is a CDN Provider Identifier as defined in Section 4.6.

Mandatory: Yes.

5.1.1.2. Trigger Status Resource

Trigger Status Resources MUST use a MIME Media Type of 'application/cdni; ptype=ci-trigger-status'.

A Trigger Status Resource is encoded as a JSON object containing the following name/value pairs.

Name: trigger

Description: The Trigger Specification posted in the body of the CI/T Command. Note that this need not be a byte-for-byte copy. For example, in the JSON representation the dCDN may re-serialise the information differently.

Value: A Trigger Specification, as defined in Section 5.2.1.

Mandatory: Yes

Name: ctime

Description: Time at which the CI/T Command was received by the dCDN. Time is determined by the dCDN, there is no requirement to synchronise clocks between interconnected CDNs.

Value: Absolute Time, as defined in Section 5.2.5.

Mandatory: Yes

Name: mtime

Description: Time at which the Trigger Status Resource was last modified. Time is determined by the dCDN, there is no requirement to synchronise clocks between interconnected CDNs.

Value: Absolute Time, as defined in Section 5.2.5.

Mandatory: Yes

Name: etime

Description: Estimate of the time at which the dCDN expects to complete the activity. Time is determined by the dCDN, there is no requirement to synchronise clocks between interconnected CDNs.

Value: Absolute Time, as defined in Section 5.2.5.

Mandatory: No

Name: status

Description: Current status of the triggered activity.

Value: Trigger Status, as defined in Section 5.2.3.

Mandatory: Yes

Name: errors

Description: Descriptions of errors that have occurred while processing a Trigger Command.

Value: An array of Error Description, as defined in Section 5.2.6. An empty array is allowed, and equivalent to omitting "errors" from the object.

Mandatory: No.

5.1.3. Trigger Collection

Trigger Collections MUST use a MIME Media Type of 'application/cdni; ptype=ci-trigger-collection'.

A Trigger Collection is encoded as a JSON object containing the following name/value pairs.

Name: triggers

Description: Links to Trigger Status Resources in the collection.

Value: A JSON array of zero or more URLs, represented as JSON strings.

Mandatory: Yes

Name: staleresourcetime

Description: The length of time for which the dCDN guarantees to keep a completed Trigger Status Resource. After this time, the dCDN SHOULD delete the Trigger Status Resource and all references to it from collections.

Value: A JSON number, which must be a positive integer, representing time in seconds.

Mandatory: Yes, in the collection of all Trigger Status Resources if the dCDN deletes stale entries. If the property is present in the filtered collections, it MUST have the same value as in the collection of all Trigger Status Resources.

Names: coll-all, coll-pending, coll-active, coll-complete, coll-failed

Description: Link to a Trigger Collection.

Value: A URL represented as a JSON string.

Mandatory: Links to all of the filtered collections are mandatory in the collection of all Trigger Status Resources, if the dCDN implements the filtered collections. Otherwise, optional.

Name: cdn-id

Description: The CDN Provider Identifier of the dCDN.

Value: A JSON string, the dCDN's CDN Provider Identifier, as defined in Section 4.6.

Mandatory: Only in the collection of all Trigger Status Resources, if the dCDN implements the filtered collections. Optional in the filtered collections (the uCDN can always find the dCDN's cdn-id in the collection of all Trigger Status Resources, but the dCDN can choose to repeat that information in its implementation of filtered collections).

5.2. Properties of CI/T Objects

This section defines the values that can appear in the top level objects described in Section 5.1, and their encodings.

5.2.1. Trigger Specification

A Trigger Collection is encoded as a JSON object containing the following name/value pairs.

An unrecognised name/value pair in the Trigger Specification object contained in a CI/T Command SHOULD be preserved in the Trigger Specification of any Trigger Status Resource it creates.

Name: type

Description: This property defines the type of the CI/T Trigger Command.

Value: Trigger Type, as defined in Section 5.2.2.

Mandatory: Yes

Name: metadata.urls

Description: The uCDN URLs of the metadata the CI/T Trigger Command applies to.

Value: A JSON array of URLs represented as JSON strings.

Mandatory: No, but at least one of 'metadata.*' or 'content.*' MUST be present and non-empty.

Name: content.urls

Description: URLs of content the CI/T Trigger Command applies to, see Section 4.8.

Value: A JSON array of URLs represented as JSON strings.

Mandatory: No, but at least one of 'metadata.*' or 'content.*' MUST be present and non-empty.

Name: content.ccid

Description: The Content Collection Identifier of content the trigger applies to. The 'ccid' is a grouping of content, as defined by [I-D.ietf-cdni-metadata].

Value: A JSON array of strings, where each string is a Content Collection Identifier.

Mandatory: No, but at least one of 'metadata.*' or 'content.*' MUST be present and non-empty.

Name: metadata.patterns

Description: The metadata the trigger applies to.

Value: A JSON array of Pattern Match, as defined in Section 5.2.4.

Mandatory: No, but at least one of 'metadata.*' or 'content.*' MUST be present and non-empty, and metadata.patterns MUST NOT be present if the TriggerType is Preposition.

Name: content.patterns

Description: The content data the trigger applies to.

Value: A JSON array of Pattern Match, as defined in Section 5.2.4.

Mandatory: No, but at least one of 'metadata.*' or 'content.*' MUST be present and non-empty, and content.patterns MUST NOT be present if the TriggerType is Preposition.

5.2.2. Trigger Type

Trigger Type is used in a Trigger Specification to describe trigger action.

All trigger types MUST be registered in the IANA CI/T Trigger Types registry (see Section 7.2).

A dCDN receiving a request containing a trigger type it does not recognise or does not support MUST reject the request by creating a Trigger Status Resource with status to "failed" and the "errors" array containing an Error Description with error "eunsupported".

The following trigger types are defined by this document:

JSON String	Description
preposition	A request for the dCDN to acquire metadata or content.
invalidate	A request for the dCDN to invalidate metadata or content. After servicing this request the dCDN will not use the specified data without first re-validating it using, for example, an "If-None-Match" HTTP request. The dCDN need not erase the associated data.
purge	A request for the dCDN to erase metadata or content. After servicing the request, the specified data MUST NOT be held on the dCDN (the dCDN should re-acquire the metadata or content from uCDN if it needs it).

5.2.3. Trigger Status

This describes the current status of a Trigger. It MUST be one of the JSON strings in the following table:

JSON String	Description
pending	The CI/T Trigger Command has not yet been acted upon.
active	The CI/T Trigger Command is currently being acted upon.
complete processed	The CI/T Trigger Command completed successfully. The CI/T Trigger Command has been accepted and no further status update will be made (can be used in cases where completion cannot be confirmed).
failed	The CI/T Trigger Command could not be completed.
cancelling	Processing of the CI/T Trigger Command is still in progress, but the CI/T Trigger Command has been cancelled by the uCDN.
cancelled	The CI/T Trigger Command was cancelled by the uCDN.

5.2.4. PatternMatch

A Pattern Match consists of a string pattern to match against a URI, and flags describing the type of match.

It is encoded as a JSON object with the following name/value pairs:

Name: pattern

Description: A pattern for URI matching.

Value: A JSON string representing the pattern. The pattern may contain the wildcards * and ?, where * matches any sequence of characters (including the empty string) and ? matches exactly one character. The three literals "\", "*" and "?" MUST be escaped as "\\ ", "*" and "\\?".

Mandatory: Yes.

Name: case-sensitive

Description: Flag indicating whether or not case-sensitive matching should be used.

Value: One of the JSON values 'true' (the matching is case-sensitive) or 'false' (the matching is case-insensitive).

Mandatory: No, default is case-insensitive match.

Name: match-query-string

Description: Flag indicating whether or not the query part of the URI should be included in the pattern match.

Value: One of the JSON values 'true' (the full URI including the query part should be compared against the given pattern), or 'false' (the query part of the URI should be dropped before comparison with the given pattern).

Mandatory: No, default is 'false', the query part of the URI should be dropped before comparison with the given pattern.

Example of case-sensitive prefix match against
"https://www.example.com/trailers/":

```
{
  "pattern": "https://www.example.com/trailers/*",
  "case-sensitive": true
}
```

5.2.5. Absolute Time

A JSON number, seconds since the UNIX epoch, 00:00:00 UTC on 1 January 1970.

5.2.6. Error Description

An Error Description is used to report failure of a CI/T Command, or in the activity it triggered. It is encoded as a JSON object with the following name/value pairs:

Name: error

Value: Error Code, as defined in Section 5.2.7.

Mandatory: Yes.

Names: metadata.urls, content.urls, metadata.patterns,
content.patterns

Description: Metadata and content references copied from the Trigger Specification. Only those URLs and patterns to which the error applies are included in each property, but those URLs and patterns MUST be exactly as they appear in the request, the dCDN MUST NOT generalise the URLs. (For example, if the uCDN requests prepositioning of URLs "https://content.example.com/a" and "https://content.example.com/b", the dCDN must not

generalise its error report to Pattern
"https://content.example.com/*".)

Value: A JSON array of JSON strings, where each string is copied from a 'content.*' or 'metadata.*' value in the corresponding Trigger Specification.

Mandatory: At least one of these name/value pairs is mandatory in each Error Description object.

Name: description

Description: A human-readable description of the error.

Value: A JSON string, the human-readable description.

Mandatory: No.

5.2.7. Error Code

This type is used by the dCDN to report failures in trigger processing. All error codes MUST be registered in the IANA CI/T Error Codes registry (see Section 7.3). Unknown error codes MUST be treated as fatal errors, and the request MUST NOT be automatically retried without modification.

The following error codes are defined by this document, and MUST be supported by an implementation of the CI/T interface.

Error Code	Description
emeta	The dCDN was unable to acquire metadata required to fulfil the request.
econtent	The dCDN was unable to acquire content (CT/T preposition commands only).
eperm	The uCDN does not have permission to issue the CI/T Command (for example, the data is owned by another CDN).
ereject	The dCDN is not willing to fulfil the CI/T Command (for example, a preposition request for content at a time when the dCDN would not accept Request Routing requests from the uCDN).
ecdn	An internal error in the dCDN or one of its downstream CDNs.
ecancelled	The uCDN cancelled the request.
eunsupported	The Trigger Specification contained a "type" that is not supported by the dCDN. No action was taken by the dCDN other than to create a Trigger Status Resource in state "failed".

6. Examples

The following sections provide examples of different CI/T objects encoded as JSON.

Discovery of the triggers interface is out of scope of this document. In an implementation, all CI/T URLs are under the control of the dCDN. The uCDN MUST NOT attempt to ascribe any meaning to individual elements of the path.

In examples in this section, the URL 'https://dcdn.example.com/triggers' is used as the location of the collection of all Trigger Status Resources, and the CDN Provider Id of uCDN is "AS64496:1".

6.1. Creating Triggers

Examples of the uCDN triggering activity in the dCDN:

6.1.1. Preposition

An example of a CI/T preposition command, a POST to the collection of all Trigger Status Resources.

Note that "metadata.patterns" and "content.patterns" are not allowed in a preposition Trigger Specification.

REQUEST:

```
POST /triggers HTTP/1.1
User-Agent: example-user-agent/0.1
Host: dcdn.example.com
Accept: */*
Content-Type: application/cdni; ptype=ci-trigger-command
Content-Length: 352

{
  "trigger" : {
    "type": "preposition",

    "metadata.urls" : [ "https://metadata.example.com/a/b/c" ],
    "content.urls" : [
      "https://www.example.com/a/b/c/1",
      "https://www.example.com/a/b/c/2",
      "https://www.example.com/a/b/c/3",
      "https://www.example.com/a/b/c/4"
    ]
  },
  "cdn-path" : [ "AS64496:1" ]
}
```

RESPONSE:

```
HTTP/1.1 201 Created
Date: Wed, 04 May 2016 08:48:10 GMT
Content-Length: 467
Content-Type: application/cdni; ptype=ci-trigger-status
Location: https://dcdn.example.com/triggers/0
Server: example-server/0.1

{
  "ctime": 1462351690,
  "etime": 1462351698,
  "mtime": 1462351690,
  "status": "pending",
  "trigger": {
    "content.urls": [
      "https://www.example.com/a/b/c/1",
      "https://www.example.com/a/b/c/2",
      "https://www.example.com/a/b/c/3",
      "https://www.example.com/a/b/c/4"
    ],
    "metadata.urls": [
      "https://metadata.example.com/a/b/c"
    ]
  },
}
```

```
        "type": "preposition"
      }
    }
```

6.1.1.2. Invalidate

An example of a CI/T invalidate command, another POST to the collection of all Trigger Status Resources. This instructs the dCDN to re-validate the content at "https://www.example.com/a/index.html", as well as any metadata and content whose URLs are prefixed by "https://metadata.example.com/a/b/" using case-insensitive matching, and "https://www.example.com/a/b/" respectively, using case-sensitive matching.

REQUEST:

```
POST /triggers HTTP/1.1
User-Agent: example-user-agent/0.1
Host: dcdn.example.com
Accept: */*
Content-Type: application/cdni; ptype=ci-trigger-command
Content-Length: 387

{
  "trigger" : {
    "type": "invalidate",

    "metadata.patterns" : [
      { "pattern" : "https://metadata.example.com/a/b/*" }
    ],

    "content.urls" : [ "https://www.example.com/a/index.html" ],
    "content.patterns" : [
      { "pattern" : "https://www.example.com/a/b/*",
        "case-sensitive" : true
      }
    ]
  },
  "cdn-path" : [ "AS64496:1" ]
}
```

RESPONSE:

```
HTTP/1.1 201 Created
Date: Wed, 04 May 2016 08:48:11 GMT
Content-Length: 545
Content-Type: application/cdni; ptype=ci-trigger-status
Location: https://dcdn.example.com/triggers/1
```

Server: example-server/0.1

```
{
  "ctime": 1462351691,
  "etime": 1462351699,
  "mtime": 1462351691,
  "status": "pending",
  "trigger": {
    "content.patterns": [
      {
        "case-sensitive": true,
        "pattern": "https://www.example.com/a/b/*"
      }
    ],
    "content.urls": [
      "https://www.example.com/a/index.html"
    ],
    "metadata.patterns": [
      {
        "pattern": "https://metadata.example.com/a/b/*"
      }
    ],
    "type": "invalidate"
  }
}
```

6.2. Examining Trigger Status

Once Trigger Status Resources have been created, the uCDN can check their status as shown in these examples.

6.2.1. Collection of All Triggers

The uCDN can fetch the collection of all Trigger Status Resources it has created that have not yet been deleted or removed as expired. After creation of the "preposition" and "invalidate" triggers shown above, this collection might look as follows:

REQUEST:

```
GET /triggers HTTP/1.1
User-Agent: example-user-agent/0.1
Host: dcdn.example.com
Accept: */*
```

RESPONSE:

```
HTTP/1.1 200 OK
Content-Length: 341
Expires: Wed, 04 May 2016 08:49:11 GMT
Server: example-server/0.1
ETag: "-936094426920308378"
Cache-Control: max-age=60
Date: Wed, 04 May 2016 08:48:11 GMT
Content-Type: application/cdni; ptype=ci-trigger-collection
```

```
{
  "cdn-id": "AS64496:0",
  "coll-active": "/triggers/active",
  "coll-complete": "/triggers/complete",
  "coll-failed": "/triggers/failed",
  "coll-pending": "/triggers/pending",
  "staleresourcetime": 86400,
  "triggers": [
    "https://dcdn.example.com/triggers/0",
    "https://dcdn.example.com/triggers/1"
  ]
}
```

6.2.2. Filtered Collections of Trigger Status Resources

The filtered collections are also available to the uCDN. Before the dCDN starts processing the two CI/T Trigger Commands shown above, both will appear in the collection of Pending Triggers, for example:

REQUEST:

```
GET /triggers/pending HTTP/1.1
User-Agent: example-user-agent/0.1
Host: dcdn.example.com
Accept: */*
```

RESPONSE:

```
HTTP/1.1 200 OK
Content-Length: 152
Expires: Wed, 04 May 2016 08:49:11 GMT
Server: example-server/0.1
ETag: "4331492443626270781"
Cache-Control: max-age=60
Date: Wed, 04 May 2016 08:48:11 GMT
Content-Type: application/cdni; ptype=ci-trigger-collection
```

```
{
  "staleresourcetime": 86400,
  "triggers": [
    "https://dcdn.example.com/triggers/0",
    "https://dcdn.example.com/triggers/1"
  ]
}
```

At this point, if no other Trigger Status Resources had been created, the other filtered views would be empty. For example:

REQUEST:

```
GET /triggers/complete HTTP/1.1
User-Agent: example-user-agent/0.1
Host: dcdn.example.com
Accept: */*
```

RESPONSE:

```
HTTP/1.1 200 OK
Content-Length: 54
Expires: Wed, 04 May 2016 08:49:11 GMT
Server: example-server/0.1
ETag: "7958041393922269003"
Cache-Control: max-age=60
Date: Wed, 04 May 2016 08:48:11 GMT
Content-Type: application/cdni; ptype=ci-trigger-collection

{
  "staleresourcetime": 86400,
  "triggers": []
}
```

6.2.3. Individual Trigger Status Resources

The Trigger Status Resources can also be examined for detail about individual CI/T Trigger Commands. For example, for the CI/T "preposition" and "invalidate" commands from previous examples:

REQUEST:

```
GET /triggers/0 HTTP/1.1
User-Agent: example-user-agent/0.1
Host: dcdn.example.com
Accept: */*
```

RESPONSE:

```
HTTP/1.1 200 OK
Content-Length: 467
Expires: Wed, 04 May 2016 08:49:10 GMT
Server: example-server/0.1
ETag: "6990548174277557683"
Cache-Control: max-age=60
Date: Wed, 04 May 2016 08:48:10 GMT
Content-Type: application/cdni; ptype=ci-trigger-status
```

```
{
  "ctime": 1462351690,
  "etime": 1462351698,
  "mtime": 1462351690,
  "status": "pending",
  "trigger": {
    "content.urls": [
      "https://www.example.com/a/b/c/1",
      "https://www.example.com/a/b/c/2",
      "https://www.example.com/a/b/c/3",
      "https://www.example.com/a/b/c/4"
    ],
    "metadata.urls": [
      "https://metadata.example.com/a/b/c"
    ],
    "type": "preposition"
  }
}
```

REQUEST:

```
GET /triggers/1 HTTP/1.1
User-Agent: example-user-agent/0.1
Host: dcdn.example.com
Accept: */*
```

RESPONSE:

```
HTTP/1.1 200 OK
Content-Length: 545
Expires: Wed, 04 May 2016 08:49:11 GMT
Server: example-server/0.1
ETag: "-554385204989405469"
Cache-Control: max-age=60
Date: Wed, 04 May 2016 08:48:11 GMT
Content-Type: application/cdni; ptype=ci-trigger-status
```

```
{
  "ctime": 1462351691,
  "etime": 1462351699,
  "mtime": 1462351691,
  "status": "pending",
  "trigger": {
    "content.patterns": [
      {
        "case-sensitive": true,
        "pattern": "https://www.example.com/a/b/*"
      }
    ],
    "content.urls": [
      "https://www.example.com/a/index.html"
    ],
    "metadata.patterns": [
      {
        "pattern": "https://metadata.example.com/a/b/*"
      }
    ],
    "type": "invalidate"
  }
}
```

6.2.4. Polling for Change

The uCDN SHOULD use the Entity Tags of collections or Trigger Status Resources when polling for change in status, as shown in the following examples:

REQUEST:

```
GET /triggers/pending HTTP/1.1
User-Agent: example-user-agent/0.1
Host: dcdn.example.com
Accept: */*
If-None-Match: "4331492443626270781"
```

RESPONSE:

```
HTTP/1.1 304 Not Modified
Content-Length: 0
Expires: Wed, 04 May 2016 08:49:11 GMT
Server: example-server/0.1
ETag: "4331492443626270781"
Cache-Control: max-age=60
Date: Wed, 04 May 2016 08:48:11 GMT
Content-Type: application/cdni; ptype=ci-trigger-collection
```

REQUEST:

```
GET /triggers/0 HTTP/1.1
User-Agent: example-user-agent/0.1
Host: dcdn.example.com
Accept: */*
If-None-Match: "6990548174277557683"
```

RESPONSE:

```
HTTP/1.1 304 Not Modified
Content-Length: 0
Expires: Wed, 04 May 2016 08:49:10 GMT
Server: example-server/0.1
ETag: "6990548174277557683"
Cache-Control: max-age=60
Date: Wed, 04 May 2016 08:48:10 GMT
Content-Type: application/cdni; ptype=ci-trigger-status
```

When the CI/T Trigger Command is complete, the contents of the filtered collections will be updated along with their Entity Tags. For example, when the two example CI/T Trigger Commands are complete, the collections of pending and complete Trigger Status Resources might look like:

REQUEST:

```
GET /triggers/pending HTTP/1.1
User-Agent: example-user-agent/0.1
Host: dcdn.example.com
Accept: */*
```

RESPONSE:

```
HTTP/1.1 200 OK
Content-Length: 54
Expires: Wed, 04 May 2016 08:49:15 GMT
Server: example-server/0.1
ETag: "1337503181677633762"
Cache-Control: max-age=60
Date: Wed, 04 May 2016 08:48:15 GMT
Content-Type: application/cdni; ptype=ci-trigger-collection

{
  "staleresourcetime": 86400,
  "triggers": []
}
```

REQUEST:

```
GET /triggers/complete HTTP/1.1
User-Agent: example-user-agent/0.1
Host: dcdn.example.com
Accept: */*
```

RESPONSE:

```
HTTP/1.1 200 OK
Content-Length: 152
Expires: Wed, 04 May 2016 08:49:22 GMT
Server: example-server/0.1
ETag: "4481489539378529796"
Cache-Control: max-age=60
Date: Wed, 04 May 2016 08:48:22 GMT
Content-Type: application/cdni; ptype=ci-trigger-collection
```

```
{
  "staleresourcetime": 86400,
  "triggers": [
    "https://dcdn.example.com/triggers/0",
    "https://dcdn.example.com/triggers/1"
  ]
}
```

6.2.5. Deleting Trigger Status Resources

The uCDN can delete completed and failed Trigger Status Resources to reduce the size of the collections, as described in Section 4.4. For example, to delete the "preposition" request from earlier examples:

REQUEST:

```
DELETE /triggers/0 HTTP/1.1
User-Agent: example-user-agent/0.1
Host: dcdn.example.com
Accept: */*
```

RESPONSE:

```
HTTP/1.1 204 No Content
Date: Wed, 04 May 2016 08:48:22 GMT
Content-Length: 0
Content-Type: text/html; charset=UTF-8
Server: example-server/0.1
```

This would, for example, cause the collection of completed Trigger Status Resources shown in the example above to be updated to:

REQUEST:

```
GET /triggers/complete HTTP/1.1
User-Agent: example-user-agent/0.1
Host: dcdn.example.com
Accept: */*
```

RESPONSE:

```
HTTP/1.1 200 OK
Content-Length: 105
Expires: Wed, 04 May 2016 08:49:22 GMT
Server: example-server/0.1
ETag: "-6938620031669085677"
Cache-Control: max-age=60
Date: Wed, 04 May 2016 08:48:22 GMT
Content-Type: application/cdni; ptype=ci-trigger-collection
```

```
{
  "staleresourcetime": 86400,
  "triggers": [
    "https://dcdn.example.com/triggers/1"
  ]
}
```

6.2.6. Error Reporting

In this example the uCDN has requested prepositioning of "https://newsite.example.com/index.html", but the dCDN was unable to locate metadata for that site:

REQUEST:

```
GET /triggers/2 HTTP/1.1
User-Agent: example-user-agent/0.1
Host: dcdn.example.com
Accept: */*
```

RESPONSE:

```
HTTP/1.1 200 OK
Content-Length: 486
Expires: Wed, 04 May 2016 08:49:26 GMT
Server: example-server/0.1
ETag: "5182824839919043757"
Cache-Control: max-age=60
Date: Wed, 04 May 2016 08:48:26 GMT
Content-Type: application/cdni; ptype=ci-trigger-status
```

```
{
  "ctime": 1462351702,
  "errors": [
    {
      "content.urls": [
        "https://newsite.example.com/index.html"
      ],
      "description": "newsite.example.com not in HostIndex",
      "error": "emeta"
    }
  ],
  "etime": 1462351710,
  "mtime": 1462351706,
  "status": "active",
  "trigger": {
    "content.urls": [
      "https://newsite.example.com/index.html"
    ],
    "type": "preposition"
  }
}
```

7. IANA Considerations

[RFC Editor Note: Please replace references to [RFCthis] in this section with this document's RFC number before publication.]

7.1. CDNI Payload Type Parameter Registrations

The IANA is requested to register the following new Payload Types in the CDNI Payload Type Parameter registry defined by [RFC7736], for use with the 'application/cdni' MIME media type.

Payload Type	Specification
ci-trigger-command	[RFCthis]
ci-trigger-status	[RFCthis]
ci-trigger-collection	[RFCthis]

7.2. CDNI CI/T Trigger Types Registry

The IANA is requested to create a new "CDNI CI/T Trigger Types" subregistry under the "Content Delivery Networks Interconnection (CDNI) Parameters" registry.

Additions to the CDNI CI/T Error Code Registry will be made via "RFC Required" as defined in [RFC5226].

The initial contents of the CDNI CI/T Trigger Types Registry comprise the names and descriptions listed in section Section 5.2.2 of this document, with this document acting as the specification.

7.3. CDNI CI/T Error Codes Registry

The IANA is requested to create a new "CDNI CI/T Error Codes" subregistry under the "Content Delivery Networks Interconnection (CDNI) Parameters" registry.

Additions to the CDNI CI/T Error Codes Registry will be made via "Specification Required" as defined in [RFC5226]. The Designated Expert will verify that new error code registrations do not duplicate existing error code definitions (in name or functionality), prevent gratuitous additions to the namespace, and prevent any additions to the namespace that would impair the interoperability of CDNI implementations.

The initial contents of the CDNI CI/T Error Codes Registry comprise the names and descriptions of the Error Codes listed in Section 5.2.7 of this document, with this document acting as the specification.

8. Security Considerations

The CI/T interface provides a mechanism to allow a uCDN to generate requests into the dCDN and to inspect its own CI/T requests and their current state. The CI/T interface does not allow access to or modification of the uCDN or dCDN metadata relating to content delivery, or to the content itself. It can only control the presence of that metadata in the dCDN, and the processing work and network utilisation involved in ensuring that presence.

By examining pre-positioning requests to a dCDN, and correctly interpreting content and metadata URLs, an attacker could learn the uCDN or content owner's predictions for future content popularity. By examining invalidate or purge requests, an attacker could learn about changes in the content owner's catalogue.

By injecting CI/T commands an attacker, or a misbehaving uCDN, would generate work in the dCDN and uCDN as they process those requests. And so would a man in the middle attacker modifying valid CI/T commands generated by the uCDN. In both cases, that would decrease the dCDN caching efficiency by causing it to unnecessarily acquire or re-acquire content metadata and/or content.

A dCDN implementation of CI/T MUST restrict the actions of a uCDN to the data corresponding to that uCDN. Failure to do so would allow uCDNs to detrimentally affect each other's efficiency by generating unnecessary acquisition or re-acquisition load.

An origin that chooses to delegate its delivery to a CDN is trusting that CDN to deliver content on its behalf, CDN-interconnection is an extension of that trust to downstream CDNs. That trust relationship is a commercial arrangement, outside the scope of the CDNi protocols. So, while a malicious CDN could deliberately generate load on a dCDN using the CI/T, the protocol does not otherwise attempt to address malicious behaviour between interconnected CDNs.

8.1. Authentication, Authorization, Confidentiality, Integrity Protection

A CI/T implementation MUST support TLS transport for HTTP (https) as per [RFC2818] and [RFC7230].

TLS MUST be used by the server-side (dCDN) and the client-side (uCDN) of the CI/T interface, including authentication of the remote end, unless alternate methods are used for ensuring the security of the information in the CI/T interface requests and responses (such as setting up an IPsec tunnel between the two CDNs or using a physically

secured internal network between two CDNs that are owned by the same corporate entity).

The use of TLS for transport of the CI/T interface allows:

- o The dCDN and the uCDN to authenticate each other using TLS client auth and TLS server auth.

And, once they have mutually authenticated each other, it allows:

- o The dCDN and the uCDN to authorize each other (to ensure they are receiving CI/T Commands from, or reporting status to, an authorized CDN).
- o CDNI commands and responses to be transmitted with confidentiality.
- o Protection of the integrity of CDNI commands and responses.

When TLS is used, the general TLS usage guidance in [RFC7525] MUST be followed.

The mechanisms for access control are dCDN-specific, not standardised as part of this CI/T specification.

HTTP requests that attempt to access or operate on CI/T data belonging to another CDN MUST be rejected using, for example, HTTP "403 Forbidden" or "404 Not Found". This is intended to prevent unauthorised users from generating unnecessary load in dCDN or uCDN due to revalidation, reacquisition, or unnecessary acquisition.

When deploying a network of interconnected CDNs, the possible inefficiencies related to the "diamond" configuration discussed in Section 2.2.1 should be considered.

8.2. Denial of Service

This document does not define a specific mechanism to protect against Denial of Service (DoS) attacks on the CI/T. However, CI/T endpoints can be protected against DoS attacks through the use of TLS transport and/or via mechanisms outside the scope of the CI/T interface, such as firewalling or use of Virtual Private Networks (VPNs).

Depending on the implementation, triggered activity may consume significant processing and bandwidth in the dCDN. A malicious or faulty uCDN could use this to generate unnecessary load in the dCDN. The dCDN should consider mechanisms to avoid overload, for example by

rate-limiting acceptance or processing of CI/T Commands, or batching up its processing.

8.3. Privacy

The CI/T protocol does not carry any information about individual End Users of a CDN, there are no privacy concerns for End Users.

The CI/T protocol does carry information which could be considered commercially sensitive by CDN operators and content owners. The use of mutually authenticated TLS to establish a secure session for the transport of CI/T data, as discussed in Section 8.1, provides confidentiality while the CI/T data is in transit, and prevents parties other than the authorised dCDN from gaining access to that data. The dCDN MUST ensure that it only exposes CI/T data related to a uCDN to clients it has authenticated as belonging to that uCDN.

9. Acknowledgements

The authors thank Kevin Ma for his input, and Carsten Bormann for his review and formalization of the JSON data.

10. References

10.1. Normative References

- [I-D.ietf-cdni-metadata]
Niven-Jenkins, B., Murray, R., Caulfield, M., and K. Ma,
"CDN Interconnection Metadata", draft-ietf-cdni-
metadata-16 (work in progress), April 2016.
- [RFC1930] Hawkinson, J. and T. Bates, "Guidelines for creation,
selection, and registration of an Autonomous System (AS)",
BCP 6, RFC 1930, DOI 10.17487/RFC1930, March 1996,
<<http://www.rfc-editor.org/info/rfc1930>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2818] Rescorla, E., "HTTP Over TLS", RFC 2818, May 2000.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an
IANA Considerations Section in RFCs", BCP 26, RFC 5226,
DOI 10.17487/RFC5226, May 2008,
<<http://www.rfc-editor.org/info/rfc5226>>.

- [RFC6707] Niven-Jenkins, B., Le Faucheur, F., and N. Bitar, "Content Distribution Network Interconnection (CDNI) Problem Statement", RFC 6707, September 2012.
- [RFC7159] Bray, T., "The JavaScript Object Notation (JSON) Data Interchange Format", RFC 7159, March 2014.
- [RFC7230] Fielding, R. and J. Reschke, "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", RFC 7230, June 2014.
- [RFC7231] Fielding, R. and J. Reschke, "Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content", RFC 7231, June 2014.
- [RFC7232] Fielding, R. and J. Reschke, "Hypertext Transfer Protocol (HTTP/1.1): Conditional Requests", RFC 7232, June 2014.
- [RFC7525] Sheffer, Y., Holz, R., and P. Saint-Andre, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", BCP 195, RFC 7525, May 2015.

10.2. Informative References

- [I-D.greevenbosch-appsawg-cbor-cddl]
Vigano, C. and H. Birkholz, "CBOR data definition language (CDDL): a notational convention to express CBOR data structures", draft-greevenbosch-appsawg-cbor-cddl-08 (work in progress), March 2016.
- [I-D.ietf-cdni-redirection]
Niven-Jenkins, B. and R. Brandenburg, "Request Routing Redirection interface for CDN Interconnection", draft-ietf-cdni-redirection-18 (work in progress), April 2016.
- [RFC7336] Peterson, L., Davie, B., and R. van Brandenburg, "Framework for Content Distribution Network Interconnection (CDNI)", RFC 7336, August 2014.
- [RFC7337] Leung, K. and Y. Lee, "Content Distribution Network Interconnection (CDNI) Requirements", RFC 7337, August 2014.
- [RFC7736] Ma, K., "Content Delivery Network Interconnection (CDNI) Media Type Registration", RFC 7736, DOI 10.17487/RFC7736, December 2015, <<http://www.rfc-editor.org/info/rfc7736>>.

Appendix A. Formalization of the JSON Data

This appendix is non-normative.

The JSON data described in this document has been formalised using CDDL [I-D.greevenbosch-appsawg-cbor-cddl] as follows:

CIT-object = CIT-command / Trigger-Status-Resource / Trigger-Collection

CIT-command ; use media type application/cdni; ptype=ci-trigger-command
= {
 ? trigger: Triggerspec
 ? cancel: [* URI]
 cdn-path: [* Cdn-PID]
}

Trigger-Status-Resource ; application/cdni; ptype=ci-trigger-status
= {
 trigger: Triggerspec
 ctime: Absolute-Time
 mtime: Absolute-Time
 ? etime: Absolute-Time
 status: Trigger-Status
 ? errors: [* Error-Description]
}

Trigger-Collection ; application/cdni; ptype=ci-trigger-collection
= {
 triggers: [* URI]
 ? staleresourcetime: int ; time in seconds
 ? coll-all: URI
 ? coll-pending: URI
 ? coll-active: URI
 ? coll-complete: URI
 ? coll-failed: URI
 ? cdn-id: Cdn-PID
}

Triggerspec = { ; 5.2.1
 type: Trigger-Type
 ? metadata.urls: [* URI]
 ? content.urls: [* URI]
 ? content.ccid: [* Ccid]
 ? metadata.patterns: [* Pattern-Match]
 ? content.patterns: [* Pattern-Match]
}

Trigger-Type = "preposition" / "invalidate" / "purge" ; 5.2.2

Trigger-Status = "pending" / "active" / "complete" / "processed"
/ "failed" / "cancelling" / "cancelled" ; 5.2.3

Pattern-Match = { ; 5.2.4
 pattern: tstr
 ? case-sensitive: bool
 ? match-query-string: bool
}

Absolute-Time = number ; seconds since UNIX epoch, 5.2.5

Error-Description = { ; 5.2.6
 error: Error-Code
 ? metadata.urls: [* URI]
 ? content.urls: [* URI]
 ? metadata.patterns: [* Pattern-Match]
 ? content.patterns: [* Pattern-Match]
 ? description: tstr
}

Error-Code = "emeta" / "econtent" / "eperm" / "ereject"
/ "ecdn" / "ecancelled" ; 5.2.7

Ccid = tstr ; see I-D.ietf-cdni-metadata

Cdn-PID = tstr .regexp "AS[0-9]+:[0-9]+"

URI = tstr

Authors' Addresses

Rob Murray
Nokia
3 Ely Road
Milton, Cambridge CB24 6DD
UK

Email: rob.murray@nokia.com

Ben Niven-Jenkins
Nokia
3 Ely Road
Milton, Cambridge CB24 6DD
UK

Email: ben.niven-jenkins@nokia.com

CDNI
Internet-Draft
Intended status: Informational
Expires: April 24, 2014

J. Seedorf
NEC
J. Peterson
Neustar
S. Previdi
Cisco
R. van Brandenburg
TNO
K. Ma
Azuki Systems, Inc.
October 21, 2013

CDNI Request Routing: Footprint and Capabilities Semantics
draft-ietf-cdni-footprint-capabilities-semantics-01

Abstract

This document tries to capture the semantics of the "Footprint and Capabilities Advertisement" part of the CDNI Request Routing interface, i.e. the desired meaning and what "Footprint and Capabilities Advertisement" is expected to offer within CDNI. The discussion in this document has the goal to facilitate the choosing of one or more suitable protocols for "Footprint and Capabilities Advertisement" within CDNI Request Routing.

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 April 24, 2014.

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 and Scope	2
2. Design Decisions for Footprint and Capabilities	4
2.1. Advertising Limited Coverage	4
2.2. Capabilities and Dynamic Data	5
2.3. Advertisement versus Queries	6
2.4. Avoiding or Handling 'cheating' dCDNs	6
2.5. Focus on Main Use Cases may Simplify Things	7
3. Main Use Case to Consider	7
4. Semantics for Footprint Advertisement	8
5. Semantics for Capabilities Advertisement	10
6. Negotiation of Support for Optional Types of Footprint/Capabilities	13
7. IANA Considerations	13
7.1. Footprint Sub-Registry	14
7.2. Protocol Sub-Registry	14
7.3. Redirection Mode Sub-Registry	14
8. Security Considerations	15
9. References	15
9.1. Normative References	15
9.2. Informative References	16
Appendix A. Acknowledgment	16
Authors' Addresses	17

1. Introduction and Scope

The CDNI working group is working on a set of protocols to enable the interconnection of multiple CDNs to a CDN federation. This CDN-federation should serve multiple purposes, as discussed in [RFC6770], for instance, to extend the reach of a given CDN to areas in the network which are not covered by this particular CDN.

The goal of this document is to achieve a clear understanding in the CDNI WG about the semantics associated with the CDNI Request Routing Footprint & Capabilities Advertisement Interface (from now on referred to as FCI), in particular the type of information a downstream CDN 'advertises' regarding its footprint and capabilities. To narrow down undecided aspects of these semantics, this document tries to establish a common understanding of what the FCI should offer and accomplish in the context of CDN Interconnection.

It is explicitly outside the scope of this document to decide on specific protocols to use for the FCI.

General assumptions in this document:

- o The CDNs participating in the CDN federation have already performed a boot strap process, i.e., they have connected to each other, either directly or indirectly, and can exchange information amongst each other.
- o The uCDN has received footprint and/or capability advertisements from a set of dCDNs. Footprint advertisement and capability advertisement need not use the same underlying protocol.
- o The upstream CDN (uCDN) receives the initial request-routing request from the endpoint requesting the resource.

The CDNI Problem Statement [RFC6707] describes footprint and capabilities advertisement as: "[enabling] a Request Routing function in an Upstream CDN to query a Request Routing function in a Downstream CDN to determine if the Downstream CDN is able (and willing) to accept the delegated Content Request". In addition, the RFC says "the CDNI Request Routing interface is also expected to enable a downstream CDN to provide to the upstream CDN (static or dynamic) information (e.g. resources, footprint, load) to facilitate selection of the downstream CDN by the upstream CDN request routing system when processing subsequent content requests from User Agents". It thus considers "resources" and "load" as capabilities to be advertised by the downstream CDN.

The range of different footprint definitions and possible capabilities is very broad. Attempting to define a comprehensive advertisement solution quickly becomes intractable. The CDNI requirements draft [I-D.ietf-cdni-requirements] lists the specific requirements for the CDNI Footprint & Capabilities Advertisement Interface in order to disambiguate footprints and capabilities with respect to CDNI. This document attempts to distill the apparent common understanding of what the terms 'footprint' and 'capabilities' mean in the context of CDNI, and detail the semantics of the

footprint advertisement mechanism and the capability advertisement mechanism.

2. Design Decisions for Footprint and Capabilities

A large part of the difficulty in discussing the FCI lies in understanding what exactly is meant when trying to define footprint in terms of "coverage" or "reachability." While the operators of CDNs pick strategic locations to situate caches, a cache with a public IPv4 address is reachable by any endpoint on the Internet unless some policy enforcement precludes the use of the cache.

Some CDNs aspire to cover the entire world, which we will henceforth call global CDNs. The footprint advertised by such a CDN in the CDNI environment would, from a coverage or reachability perspective, presumably cover all prefixes. Potentially more interesting for CDNI use cases, however, are CDNs that claim a more limited coverage, but seek to federate with other CDNs in order to create a single CDN fabric which shares resources.

Futhermore, not all capabilities need be footprint restricted. Depending upon the use case, the optimal semantics of "footprints with capability attributes" vs. "capabilities with footprint restrictions" are not clear.

The key to understanding the semantics of footprint and capability advertisement lies in understand why a dCDN would advertise a limited coverage area, and how a uCDN would use such advertisements to decide among one of several dCDNs. The following section will discuss some of the trade-offs and design decisions that need to be decided upon for the CDNI FCI.

2.1. Advertising Limited Coverage

The basic use case that would motivate a dCDN to advertise a limited coverage is that the CDN was built to cover only a particular portion of the Internet. For example, an ISP could purpose-build a CDN to serve only their own customers by situating caches in close topological proximity to high concentrations of their subscribers. The ISP knows the prefixes it has allocated to end users and thus can easily construct a list of prefixes that its caches were positioned to serve.

When such a purpose-built CDN joins a federation, however, and advertises its footprint to a uCDN, the original intended coverage of the CDN might not represent its actual value to the federation of CDNs. Consider an ISP-A and ISP-B that both field their own CDNs, which they federate through CDNI. A given user E, who is customer of

ISP-B, might happen to be topologically closest to a cache fielded by ISP-A, if E happens to live in a region where ISP-B has few customers and ISP-A has many. In this case, should ISP-A's CDN "cover" E? If ISP-B's CDN has a failure condition, should the uCDN understand that ISP-A's caches are potentially available back-ups - and if so, how does ISP-A advertise itself as a "standby" for E? What about the case where CDNs advertising to the same uCDN express overlapping coverage (for example, a federation mixing global and limited CDNs)?

The answers to these questions greatly depend on how much information we want the uCDN to use to make a selection of a dCDN. If a uCDN has three dCDNs to choose from that "cover" the IP address of user E, obviously the uCDN might be interested to know how optimal the coverage is from each of the dCDNs - coverage need not be binary, either provided or not provided. dCDNs could advertise a coverage "score," for example, and provided that they all reported scores fairly on the same scale, uCDNs could use that to make their topological optimality decision. Alternatively, dCDNs could for their footprint advertise the IP addresses of their caches rather than prefix "coverage," and let the uCDN decide for itself (based on its own topological intelligence) which dCDN has better resources to serve a given user.

In summary, the semantics of advertising footprint depend on whether such qualitative metrics for expressing footprint (such as the coverage 'score' mentioned above) should be part of the CDNI FCI, or if it should focus just on 'binary' footprint.

2.2. Capabilities and Dynamic Data

In cases where the apparent footprint of dCDNs overlaps, uCDNs might also want to rely on a host of other factors to evaluate the respective merits of dCDNs. These include facts related to the caches themselves, to the network where the cache is deployed, to the nature of the resource sought and to the administrative policies of the respective networks.

In the absence of network-layer impediments to reaching caches, the choice to limit coverage is necessarily an administrative policy. Much policy must be agreed upon before CDNs can merge into federations, including questions of membership, compensation, volumes and so on. A uCDN certainly will factor these sorts of considerations into its decision to select a dCDN, but there is probably little need for dCDNs to actually advertise them through an interface - they will be settled out of band as a precondition for federating.

Other facts about the dCDN would be expressed through the interface to the uCDN. Some capabilities of a dCDN are static, and some are highly dynamic. Expressing the total storage built into its caches, for example, changes relatively rarely, whereas the amount of storage in use at any given moment is highly volatile. Network bandwidth similarly could be expressed as either total bandwidth available to a cache, or based on the current state of the network. A cache may at one moment lack a particular resource in storage, but have it the next.

The semantics of the capabilities interface will depend on how much of the dCDN state needs to be pushed to the uCDN and qualitatively how often that information should be updated.

2.3. Advertisement versus Queries

In a federated CDN environment, each dCDN shares some of its state with the uCDN, which the uCDN uses to build a unified picture of all of the dCDNs available to it. In architectures that share detailed capability information, the uCDN could basically perform the entire request-routing intelligence down to selecting a particular cache before sending the request to the dCDN (note that within the current CDNI WG scope, such direct selection of specific caches by the uCDN is out of scope). However, when the uCDN must deal with many potential dCDNs, this approach does not scale. Especially as CDNs scale up from dozens or hundreds of caches to thousands or tens of thousands, the volume of updates to footprint and capability may become onerous.

Were the volume of updates to exceed the volumes of requests to the uCDN, it might make more sense for the uCDN to query dCDNs upon receiving requests (as is the case in the recursive redirection mode described in [I-D.ietf-cdni-framework]), instead of receiving advertisements and tracking the state of dCDNs itself. The advantage of querying dCDNs would be that much of the dynamic data that dCDNs cannot share with the uCDN would now be factored into the uCDN's decision. dCDNs need not replicate any state to the uCDN - uCDNs could effectively operate in a stateless mode.

The semantics of both footprint and capability advertisement depend on the service model here: are there cases where a synchronous query/response model would work better for the uCDN decision than a state replication model?

2.4. Avoiding or Handling 'cheating' dCDNs

In a situation where more than one dCDN is willing to serve a given end user request, it might be attractive for a dCDN to 'cheat' in the

sense that the dCDN provides inaccurate information to the uCDN in order to convince the uCDN to select it opposed to 'competing' dCDNs. It could therefore be desirable to take away the incentive for dCDNs to cheat (in information advertised) as much as possible. One option here is to make the information the dCDN advertises somehow verifiable for the uCDN. On the other hand, a cheating dCDN might be avoided or handled by the fact that there will be strong contractual agreements between a uCDN and a dCDN, so that a dCDN would risk severe penalties or legal consequences when caught cheating.

Overall, it seems that information a dCDN advertises should (in the long run) be somehow qualitatively verifiable by the uCDN, though possibly through non-real-time out-of-band audits. It is probably an overly strict requirement to mandate that such verification be possible "immediately", i.e. during the request routing process itself. If the uCDN can detect a cheating dCDN at a later stage, it should suffice for the uCDN to "de-incentivize" cheating because it would negatively affect the long-term business relationship with a particular dCDN.

2.5. Focus on Main Use Cases may Simplify Things

To narrow down semantics for "footprint" and "capabilities" in the CDNI context, it can be useful to initially focus on key use cases to be addressed by the CDNI WG that are to be envisioned the main deployments in the foreseeable future. In this regard, a main realistic use case is the existence of ISP-owned CDNs, which essentially cover a certain operator's network. At the same time, however, the possibility of overlapping footprints should not be excluded, i.e. the scenario where more than one dCDN claims it can serve a given end user request. The ISPs may also choose to federate with a fallback global CDN.

It seems reasonable to assume that in most use cases it is the uCDN that makes the decision on selecting a certain dCDN for request routing based on information the uCDN has received from this particular dCDN. It may be assumed that 'cheating' CDNs will be dealt with via means outside the scope of CDNI and that the information advertised between CDNs is accurate. In addition, excluding the use of qualitative information (e.g., cache proximity, delivery latency, cache load) to predict the quality of delivery would further simplify the use case allowing it to better focus on the basic functionality of the FCI.

3. Main Use Case to Consider

Focusing on a main use case that contains a simple (yet somewhat challenging), realistic, and generally imaginable scenario can help in narrowing down the requirements for the CDNI FCI. To this end, the following (simplified) use case can help in clarifying the semantics of footprint and capabilities for CDNI. In particular, the intention of the use case is to clarify what information needs to be exchanged on the CDNI FCI, what types of information need to be supported in a mandatory fashion (and which should be considered optional), and what types of information need to be updated with respect to a priori established CDNI contracts.

In short, one can imagine the following use case: A given uCDN has several dCDNs. It selects one dCDN for delivery protocol A and footprint 1 and another dCDN for delivery protocol B and footprint 1. The dCDN that serves delivery protocol B has a further, transitive (level-2) dCDN, that serves delivery protocol B in a subset of footprint 1 where the first-level dCDN cannot serve delivery protocol B itself. What happens if capabilities change in the transitive level-2 dCDN that might affect how the uCDN selects a level-1 dCDN (e.g. in case the level-2 dCDN cannot serve delivery protocol B anymore)? How will these changes be conveyed to the uCDN? In particular, what information does the uCDN need to be able to select a new first-level dCDN, either for all of footprint 1 or only for the subset of footprint 1 that the transitive level-2 dCDN served on behalf of the first-level dCDN?

4. Semantics for Footprint Advertisement

Roughly speaking, "footprint" can be defined as "ability and willingness to serve" by a downstream CDN. However, in addition to simple "ability and willingness to serve", the uCDN may wish to have additional information to make a dCDN selection decision, e.g., "how well" a given dCDN can actually serve a given end user request. The "ability and willingness" to serve should be distinguished from the subjective qualitative measurement of "how well" it was served. One can imagine that such additional information is implicitly associated with a given footprint, e.g. due to contractual agreements (e.g. SLAs), business relationships, or perceived dCDN quality in the past. As an alternative, such additional information could also be explicitly tagged along with the footprint.

It is reasonable to assume that a significant part of the actual footprint advertisement will happen in contractual agreements between participating CDNs, i.e. prior to the advertisement phase using the CDNI FCI. The reason for this assumption is that any contractual agreement is likely to contain specifics about the dCDN coverage (i.e. the dCDN footprint) the contractual agreement applies to. In particular, additional information to judge the delivery quality

associated with a given dCDN footprint might be defined in contractual agreements (i.e. outside of the CDNI FCI). Further, one can assume that dCDN contractual agreements about the delivery quality associated with a given footprint will probably be based on high-level aggregated statistics (i.e. not too detailed).

Given that a large part of footprint advertisement will actually happen in contractual agreements, the semantics of CDNI footprint advertisement refer to answering the following question: what exactly still needs to be advertised by the CDNI FCI? For instance, updates about temporal failures of part of a footprint can be useful information to convey via the CDNI request routing interface. Such information would provide updates on information previously agreed in contracts between the participating CDNs. In other words, the CDNI FCI is a means for a dCDN to provide changes/updates regarding a footprint it has prior agreed to serve in a contract with a uCDN.

Generally speaking, one can imagine two categories of footprint to be advertised by a dCDN:

- o Footprint could be defined based on "coverage/reachability", where coverage/reachability refers to a set of prefixes, a geographic region, or similar boundary. The dCDN claims that it can cover/reach 'end user requests coming from this footprint'.
- o Footprint could be defined based on "resources", where resources refers to surrogates/caches a dCDN claims to have (e.g., the location of surrogates/resources). The dCDN claims that 'from this footprint' it can serve incoming end user requests.

For each of these footprint types, there are capabilities associated with a given footprint, i.e. the capabilities (e.g., delivery protocol, redirection mode, metadata) supported in the coverage area for a "coverage/reachability" defined footprint, or the capabilities of resources (e.g., delivery protocol, redirection mode, metadata support) for a "resources" defined footprint.

It seems clear that "coverage/reachability" types of footprint must be supported within CDNI. The following such types of footprint are mandatory and must be supported by the CDNI FCI:

- o List of ISO Country Codes
- o List of AS numbers
- o Set of IP-prefixes

A 'set of IP-prefixes' must be able to contain full IP addresses, i.e., a /32 for IPv4 and a /128 for IPv6, and also IP prefixes with an arbitrary prefix length. There must also be support for multiple IP address versions, i.e., IPv4 and IPv6, in such a footprint.

For all of these mandatory-to-implement footprint types, footprints can be viewed as constraints for delegating requests to a dCDN: A dCDN footprint advertisement tells the uCDN the limitations for delegating a request to the dCDN. For IP prefixes or ASN(s), the footprint signals to the uCDN that it should consider the dCDN a candidate only if the IP address of the request routing source falls within the prefix set (or ASN, respectively). The CDNI specifications do not define how a given uCDN determines what address ranges are in a particular ASN. Similarly, for country codes a uCDN should only consider the dCDN a candidate if it covers the country of the request routing source. The CDNI specifications do not define how a given uCDN determines the country of the request routing source. Multiple footprint constraints are additive, i.e. the advertisement of different types of footprint narrows the dCDN candidacy cumulatively.

In addition to these mandatory "coverage/reachability" types of footprint, other optional "coverage/reachability" types of footprint or "resource" types of footprint may be defined by future specifications. To facilitate this, a clear process for specifying optional footprint types in a IANA registry must be specified. This includes the specification of the level of oversight necessary (e.g. WG decision or expert review) for adding new optional footprints to a IANA registry as well as the specification of a template regarding design choices that must be captured by new optional types of footprints.

Independent of the exact type of a footprint, a footprint might also include the connectivity of a given dCDN to other CDNs that may be able to serve content to users on behalf of that dCDN, to cover cases where there is a transitive CDN interconnection. Further, the downstream CDN must be able to express its footprint to an interested upstream CDN (uCDN) in a comprehensive form, e.g., as a complete data set containing the complete footprint. Making incremental updates, however, to express dynamic changes in state is also desirable.

5. Semantics for Capabilities Advertisement

In general, the dCDN must be able to express its general capabilities to the uCDN. These general capabilities could express if the dCDN supports a given service, for instance, HTTP delivery, RTP/RTSP delivery or RTMP. Furthermore, the dCDN must be able to express particular capabilities for the delivery in a particular footprint

area. For example, the dCDN might in general offer RTMP but not in some specific areas, either for maintenance reasons or because the caches covering this particular area cannot deliver this type of service. Hence, in certain cases footprint and capabilities are tied together and cannot be interpreted independently from each other. In such cases, i.e. where capabilities must be expressed on a per footprint basis, it may be beneficial to combine footprint and capabilities advertisement.

A high-level and very rough semantic for capabilities is thus the following: Capabilities are types of information that allow a uCDN to determine if a downstream CDN is able (and willing) to accept (and properly handle) a delegated content request. In addition, Capabilities are characterized by the fact that this information may possibly change over time based on the state of the network or caches.

At a first glance, several broad categories of capabilities seem useful to convey via an advertisement interface, however, advertising capabilities that change highly dynamically (e.g. real-time delivery performance metrics, CDN resource load, or other highly dynamically changing QoS information) should probably not be in scope for the CDNI FCI. First, out of the multitude of possible metrics and capabilities, it is hard to agree on a subset and the precise metrics to be used. Second, and perhaps more importantly, it seems not feasible to specify such highly dynamically changing capabilities and the corresponding metrics within the CDNI charter time-frame.

Useful capabilities refer to information that does not change highly dynamically and which in many cases is absolutely necessary to decide on a particular dCDN for a given end user request. For instance, if an end user request concerns the delivery of a video file with a certain protocol (e.g. RTMP), the uCDN needs to know if a given dCDN has the capability of supporting this delivery protocol.

Similar to footprint advertisement, it is reasonable to assume that a significant part of the actual (resource) capabilities advertisement will happen in contractual agreements between participating CDNs, i.e. prior to the advertisement phase using the CDNI FCI. The role of capability advertisement is hence rather to enable the dCDN to update a uCDN on changes since a contract has been set up (e.g. in case a new delivery protocol is suddenly being added to the list of supported delivery protocols of a given dCDN, or in case a certain delivery protocol is suddenly not being supported anymore due to failures). Capabilities advertisement thus refers to conveying information to a uCDN about changes/updates of certain capabilities with respect to a given contract.

Given these semantics, it needs to be decided what exact capabilities are useful and how these can be expressed. Since the details of CDNI contracts are not known at the time of this writing (and the CDNI interface should probably be agnostic to these contracts anyway), it remains to be seen what capabilities will be used to define agreements between CDNs in practice. One implication for standardization may be to initially only specify a very limited set of mandatory capabilities for advertisement and have on top of that a flexible data model that allows exchanging additional capabilities when needed. Still, agreement needs to be found on which capabilities (if any) should be mandatory among CDNs. As discussed in Section 2.5, finding the concrete answers to these questions can benefit from focusing on a small number of key use cases that are highly relevant and contain enough complexity to help in understanding what concrete capabilities are needed to facilitate CDN Interconnection.

Under the above considerations, the following capabilities seem useful as 'base' capabilities, i.e. ones that are needed in any case and therefore constitute mandatory capabilities to be supported by the CDNI FCI:

- o Delivery Protocol (e.g., HTTP vs. RTMP)
- o Acquisition Protocol (for acquiring content from a uCDN)
- o Redirection Mode (e.g., DNS Redirection vs. HTTP Redirection as discussed in [I-D.ietf-cdni-framework])
- o Capabilities related to CDNI Logging (e.g., supported logging mechanisms)
- o Capabilities related to CDNI Metadata (e.g., authorization algorithms or support for proprietary vendor metadata)

It is not feasible to enumerate all the possible options for the mandatory capabilities listed above (e.g., all the potential delivery protocols or metadata options) or anticipate all the future needs for additional capabilities. It would be unreasonable to burden the CDNI FCI specification with defining each supported capability. Instead, the CDNI FCI specification should define a generic protocol for conveying any capability information. In this respect, it seems reasonable to define a registry which initially contains the mandatory capabilities listed above, but may be extended as needs dictate. The CDNI FCI specification SHOULD define the registry (and the rules for adding new entries to the registry) for the different capability types. Each capability type MAY further have a list of valid values. The individual CDNI interface specifications which

define a given capability SHOULD define any necessary registries (and the rules for adding new entries to the registry) for the values advertised for a given capability type.

The mandatory capabilities listed above generally relate to information that is configured on a content asset or group of assets basis via CDNI metadata. The capability requirements for acquisition and delivery protocol and other mandatory metadata capabilities (e.g. authorization algorithms) are defined in [I-D.ietf-cdni-metadata].

Note: CDNI interface support for logging configuration (i.e., control interface vs. metadata interface) has not yet been decided. Once it has been decided, the corresponding CDNI interface specification should define the associated capability requirements.

6. Negotiation of Support for Optional Types of Footprint/Capabilities

The notion of optional types of footprint and capabilities implies that certain implementations may not support all kinds of footprint and capabilities. Therefore, any FCI solution protocol must define how the support for optional types of footprint/capabilities will be negotiated between a uCDN and a dCDN that use the particular FCI protocol. In particular, any FCI solution protocol needs to specify how to handle failure cases or non-supported types of footprint/capabilities.

In general, a uCDN may ignore capabilities or types of footprint it does not understand; in this case it only selects a suitable downstream CDN based on the types of capabilities and footprint it understands. Similarly, if a dCDN does not use an optional capability or footprint which is, however, supported by a uCDN, this causes no problem for the FCI functionality because the uCDN decides on the remaining capabilities/footprint information that is being conveyed by the dCDN.

7. IANA Considerations

IANA registries are to be used for mandatory and optional types of footprint and capabilities. Therefore, the mandatory types of footprint and capabilities listed in this document (see Section 5) are to be registered with IANA. In order to prevent namespace collisions for capabilities a new IANA registry is requested for the "CDNI Capabilities" namespace. The namespace shall be split into two partitions: standard and vendor defined. As with the CDNI Metadata Interface [I-D.ietf-cdni-metadata], the vendor defined namespace partition SHOULD use a namespace prefix of "ext.", while the standard namespace partition MUST NOT.

The standard namespace partition MUST conform to the "RFC Required" policy as defined in [RFC5226]. The vendor defined namespace partition should be further partitioned into vendor specific partitions with the prefix "ext.vendor_name.". The vendor defined partition SHOULD conform to the "Expert Review" policy as defined in [RFC5226]. The expert review is simply to prevent namespace hoarding. The vendor specific partitions MAY conform to the "First Come First Served" policy as defined in [RFC5226], however, vendors defining new capabilities which conflict with existing capabilities SHOULD follow the guidelines for the "Specification Required" policy as defined in [RFC5226].

The following table defines the initial capabilities for the standard partition:

capability	Specification
Delivery Protocol	RFCthis
Acquisition Protocol	RFCthis
Redirection Mode	RFCthis

Additional capabilities specific to the CDNI Metadata Interface [I-D.ietf-cdni-metadata] and the CDNI Logging Interface [I-D.ietf-cdni-metadata] SHALL be addressed separately by those documents.

7.1. Footprint Sub-Registry

The "CDNI Metadata Footprint Types" namespace defined in the CDNI Metadata Interface document [I-D.ietf-cdni-metadata] contains the supported footprint formats for use in footprint advertisement. No further IANA action is required here.

7.2. Protocol Sub-Registry

The "CDNI Metadata Protocols" namespace defined in the CDNI Metadata Interface document [I-D.ietf-cdni-metadata] contains the supported protocol values for the Delivery Protocol and Acquisition Protocol capabilities. No further IANA action is required here.

7.3. Redirection Mode Sub-Registry

The "CDNI Capabilities Redirection Modes" namespace defines the valid redirection modes that may be advertised as supported by a CDN. Additions to the Redirection Mode namespace MUST conform to the "Expert Review" policy as defined in [RFC5226]. The expert review should verify that new type definitions do not duplicate existing type definitions and prevent gratuitous additions to the namespace. For new Redirection Modes which apply to new standard protocols, it is recommended that registration requests follow the "RFC Required" policy as defined in [RFC5226].

The following table defines the initial Redirection Modes:

Redirection Mode	definition
DNS-I	Iterative DNS-based Redirection
DNS-R	Recursive DNS-based Redirection
HTTP-I	Iterative HTTP-based Redirection
HTTP-R	Recursive HTTP-based Redirection

8. Security Considerations

Security considerations will be discussed in a future version of this document.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [RFC6707] Niven-Jenkins, B., Le Faucheur, F., and N. Bitar, "Content Distribution Network Interconnection (CDNI) Problem Statement", RFC 6707, September 2012.
- [RFC6770] Bertrand, G., Stephan, E., Burbridge, T., Eardley, P., Ma, K., and G. Watson, "Use Cases for Content Delivery Network Interconnection", RFC 6770, November 2012.

9.2. Informative References

[I-D.ietf-cdni-framework]

Peterson, L. and B. Davie, "Framework for CDN Interconnection", draft-ietf-cdni-framework-06 (work in progress), October 2013.

[I-D.ietf-cdni-logging]

Faucheur, F., Bertrand, G., Oprescu, I., and R. Peterkofsky, "CDNI Logging Interface", draft-ietf-cdni-logging-08 (work in progress), October 2013.

[I-D.ietf-cdni-metadata]

Niven-Jenkins, B., Murray, R., Watson, G., Caulfield, M., Leung, K., and K. Ma, "CDN Interconnect Metadata", draft-ietf-cdni-metadata-02 (work in progress), July 2013.

[I-D.ietf-cdni-requirements]

Leung, K. and Y. Lee, "Content Distribution Network Interconnection (CDNI) Requirements", draft-ietf-cdni-requirements-10 (work in progress), September 2013.

Appendix A. Acknowledgment

Jan Seedorf is partially supported by the CHANGE project (CHANGE: Enabling Innovation in the Internet Architecture through Flexible Flow-Processing Extensions, <http://www.change-project.eu/>), a research project supported by the European Commission under its 7th Framework Program (contract no. 257422). The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the CHANGE project or the European Commission.

Jan Seedorf has been partially supported by the COAST project (Content Aware Searching, retrieval and sTreaming, <http://www.coast-fp7.eu>), a research project supported by the European Commission under its 7th Framework Program (contract no. 248036). The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the COAST project or the European Commission.

Martin Stiernerling provided initial input to this document and valuable comments to the ongoing discussions among the authors of this document. Thanks to Francois Le Faucheur and Scott Wainner for providing valuable comments and suggestions to the text.

Authors' Addresses

Jan Seedorf
NEC
Kurfuerstenanlage 36
Heidelberg 69115
Germany

Phone: +49 6221 4342 221
Fax: +49 6221 4342 155
Email: seedorf@neclab.eu

Jon Peterson
NeuStar
1800 Sutter St Suite 570
Concord CA 94520
USA

Email: jon.peterson@neustar.biz

Stefano Previdi
Cisco Systems
Via Del Serafico 200
Rome 0144
Italy

Email: sprevidi@cisco.com

Ray van Brandenburg
TNO
Brassersplein 2
Delft 2612CT
The Netherlands

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

Kevin J. Ma
Azuki Systems, Inc.
43 Nagog Park
Acton MA 01720
USA

Phone: +1 978-844-5100
Email: kevin.ma@azukisystems.com

CDNI
Internet-Draft
Intended status: Standards Track
Expires: November 21, 2016

J. Seedorf
NEC
J. Peterson
Neustar
S. Previdi
Cisco
R. van Brandenburg
TNO
K. Ma
Ericsson
May 20, 2016

CDNI Request Routing: Footprint and Capabilities Semantics
draft-ietf-cdni-footprint-capabilities-semantics-20

Abstract

This document captures the semantics of the "Footprint and Capabilities Advertisement" part of the CDNI Request Routing interface, i.e., the desired meaning of "Footprint" and "Capabilities" in the CDNI context, and what the "Footprint and Capabilities Advertisement Interface (FCI)" offers within CDNI. The document also provides guidelines for the CDNI FCI protocol. It further defines a Base Advertisement Object, the necessary registries for capabilities and footprints, and guidelines on how these registries can be extended in the future.

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 November 21, 2016.

Copyright Notice

Copyright (c) 2016 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 and Scope	3
1.1. Terminology	4
2. Design Decisions for Footprint and Capabilities	5
2.1. Advertising Limited Coverage	5
2.2. Capabilities and Dynamic Data	6
2.3. Advertisement versus Queries	7
2.4. Avoiding or Handling 'cheating' dCDNs	7
3. Focusing on Capabilities with Footprint Restrictions	8
4. Footprint and Capabilities Extension	8
5. Capability Advertisement Object	10
5.1. Base Advertisement Object	10
5.2. Encoding	11
5.3. Delivery Protocol Capability Object	11
5.3.1. Delivery Protocol Capability Object Serialization	12
5.4. Acquisition Protocol Capability Object	12
5.4.1. Acquisition Protocol Capability Object Serialization	13
5.5. Redirection Mode Capability Object	13
5.5.1. Redirection Mode Capability Object Serialization	13
5.6. CDNI Logging Capability Object	14
5.6.1. CDNI Logging Capability Object Serialization	15
5.7. CDNI Metadata Capability Object	15
5.7.1. CDNI Metadata Capability Object Serialization	16
6. IANA Considerations	17
6.1. CDNI Payload Types	17
6.1.1. CDNI FCI DeliveryProtocol Payload Type	17
6.1.2. CDNI FCI AcquisitionProtocol Payload Type	18
6.1.3. CDNI FCI RedirectionMode Payload Type	18
6.1.4. CDNI FCI Logging Payload Type	18
6.1.5. CDNI FCI Metadata Payload Type	18

6.2. Redirection Mode Registry	18
7. Security Considerations	19
8. References	20
8.1. Normative References	20
8.2. Informative References	20
Appendix A. Main Use Case to Consider	21
Appendix B. Semantics for Footprint Advertisement	22
Appendix C. Semantics for Capabilities Advertisement	24
Appendix D. Acknowledgment	25
Authors' Addresses	26

1. Introduction and Scope

The CDNI working group is working on a set of protocols to enable the interconnection of multiple CDNs. This CDN interconnection (CDNI) can serve multiple purposes, as discussed in [RFC6770], for instance, to extend the reach of a given CDN to areas in the network which are not covered by this particular CDN.

The goal of this document is to achieve a clear understanding about the semantics associated with the CDNI Request Routing Footprint & Capabilities Advertisement Interface (from now on referred to as FCI), in particular the type of information a downstream CDN (dCDN) 'advertises' regarding its footprint and capabilities. To narrow down undecided aspects of these semantics, this document tries to establish a common understanding of what the FCI needs to offer and accomplish in the context of CDNI.

It is explicitly outside the scope of this document to decide on specific protocols to use for the FCI. However, guidelines for such FCI protocols are provided.

General assumptions in this document:

- o The CDNs participating in the interconnected CDN have already performed a boot strap process, i.e., they have connected to each other, either directly or indirectly, and can exchange information amongst each other.
- o The upstream CDN (uCDN) receives footprint and/or capability advertisements from a set of dCDNs. Footprint advertisement and capability advertisement need not use the same underlying protocol.
- o The uCDN receives the initial request-routing request from the endpoint requesting the resource.

The CDNI Problem Statement [RFC6707] describes the Request Routing Interface as: "[enabling] a Request Routing function in a uCDN to query a Request Routing function in a dCDN to determine if the dCDN is able (and willing) to accept the delegated Content Request". In addition, RFC6707 says "the CDNI Request Routing interface is also expected to enable a dCDN to provide to the uCDN (static or dynamic) information (e.g., resources, footprint, load) to facilitate selection of the dCDN by the uCDN request routing system when processing subsequent content requests from User Agents". It thus considers "resources" and "load" as capabilities to be advertised by the dCDN.

The range of different footprint definitions and possible capabilities is very broad. Attempting to define a comprehensive advertisement solution quickly becomes intractable. The CDNI requirements draft [RFC7337] lists the specific requirements for the CDNI Footprint & Capabilities Advertisement Interface in order to disambiguate footprints and capabilities with respect to CDNI. This document defines a common understanding of what the terms 'footprint' and 'capabilities' mean in the context of CDNI, and details the semantics of the footprint advertisement mechanism and the capability advertisement mechanism.

1.1. Terminology

This document reuses the terminology defined in [RFC6707].

Additionally, the following terms are used throughout this document and are defined as follows:

- o **Footprint:** a description of a CDN's coverage area, i.e., the area from which client requests may originate for, and to which the CDN is willing to deliver, content. Note: There are many ways to describe a footprint, for example, by address range (e.g., IPv4/IPv6 CIDR), by network ID (e.g., ASN), by nation boundaries (e.g., country code), by GPS coordinates, etc. This document does not define or endorse the quality or suitability of any particular footprint description method; this document only defines a method for transporting known footprint descriptions in Footprint and Capabilities Advertisement messages.
- o **Capability:** a feature of a dCDN, upon which a uCDN relies on the dCDN supporting, when making delegation decisions. Support for a given feature can change over time and can be restricted to a limited portion of a dCDN's footprint. Note: There are many possible dCDN features that could be of interest to a uCDN. This document does not presume to define them all; this document describes a scheme for defining new capabilities and how to

transport them in Footprint and Capabilities Advertisement messages.

2. Design Decisions for Footprint and Capabilities

A large part of the difficulty in discussing the FCI lies in understanding what exactly is meant when trying to define footprint in terms of "coverage" or "reachability." While the operators of CDNs pick strategic locations to situate surrogates, a surrogate with a public IPv4 address is reachable by any endpoint on the Internet unless some policy enforcement precludes the use of the surrogate.

Some CDNs aspire to cover the entire world; we refer to these as global CDNs. The footprint advertised by such a CDN in the CDNI environment would, from a coverage or reachability perspective, presumably cover all prefixes. Potentially more interesting for CDNI use cases, however, are CDNs that claim a more limited coverage, but seek to interconnect with other CDNs in order to create a single CDN fabric which shares resources.

Furthermore, not all capabilities need to be footprint restricted. Depending upon the use case, the optimal semantics of "footprints with capability attributes" vs. "capabilities with footprint restrictions" are not clear.

The key to understanding the semantics of footprint and capability advertisement lies in understanding why a dCDN would advertise a limited coverage area, and how a uCDN would use such advertisements to decide among one of several dCDNs. The following section will discuss some of the trade-offs and design decisions that need to be decided upon for the CDNI FCI.

2.1. Advertising Limited Coverage

The basic use case that would motivate a dCDN to advertise a limited coverage is that the CDN was built to cover only a particular portion of the Internet. For example, an ISP could purpose-build a CDN to serve only their own customers by situating surrogates in close topological proximity to high concentrations of their subscribers. The ISP knows the prefixes it has allocated to end users and thus can easily construct a list of prefixes that its surrogates were positioned to serve.

When such a purpose-built CDN interconnects with other CDNs and advertises its footprint to a uCDN, however, the original intended coverage of the CDN might not represent its actual value to the interconnection of CDNs. Consider an ISP-A and ISP-B that both field their own CDNs, which they interconnect via CDNI. A given user E,

who is a customer of ISP-B, might happen to be topologically closer to a surrogate fielded by ISP-A, if E happens to live in a region where ISP-B has few customers and ISP-A has many. In this case, is it ISP-A's CDN that "covers" E? If ISP-B's CDN has a failure condition, is it up to the uCDN to understand that ISP-A's surrogates are potentially available as back-ups - and if so, how does ISP-A advertise itself as a "standby" for E? What about the case where CDNs advertising to the same uCDN express overlapping coverage (for example, mixing global and limited CDNs)?

The answers to these questions greatly depend on how much information the uCDN wants to use to make a selection of a dCDN. If a uCDN has three dCDNs to choose from that "cover" the IP address of user E, obviously the uCDN might be interested to know how optimal the coverage is from each of the dCDNs - coverage need not be binary, either provided or not provided. dCDNs could advertise a coverage "score," for example, and provided that they all reported scores fairly on the same scale, uCDNs could use that to make their topological optimality decision. Alternately, dCDNs could advertise the IP addresses of their surrogates rather than prefix "coverage," and let the uCDN decide for itself (based on its own topological intelligence) which dCDN has better resources to serve a given user.

In summary, the semantics of advertising footprint depend on whether such qualitative metrics for expressing footprint (such as the coverage 'score' mentioned above) are included as part of the CDNI FCI, or if the focus is just on 'binary' footprint.

2.2. Capabilities and Dynamic Data

In cases where the apparent footprints of dCDNs overlap, uCDNs might also want to rely on other factors to evaluate the respective merits of dCDNs. These include facts related to the surrogates themselves, to the network where the surrogate is deployed, to the nature of the resource sought, and to the administrative policies of the respective networks.

In the absence of network-layer impediments to reaching surrogates, the choice to limit coverage is necessarily an administrative policy. Much policy needs to be agreed upon before CDNs can interconnect, including questions of membership, compensation, volumes, and so on. A uCDN certainly will factor these sorts of considerations into its decision to select a dCDN, but there is probably little need for dCDNs to actually advertise them through an interface - they will be settled out-of-band as a precondition for interconnection.

Other facts about the dCDN would be expressed through the interface to the uCDN. Some capabilities of a dCDN are static, and some are

highly dynamic. Expressing the total storage built into its surrogates, for example, changes relatively rarely, whereas the amount of storage in use at any given moment is highly volatile. Network bandwidth similarly could be expressed as either total bandwidth available to a surrogate, or based on the current state of the network. A surrogate can at one moment lack a particular resource in storage, but have it the next.

The semantics of the capabilities interface will depend on how much of the dCDN state needs to be pushed to the uCDN and qualitatively how often that information needs to be updated.

2.3. Advertisement versus Queries

In a CDNI environment, each dCDN shares some of its state with the uCDN. The uCDN uses this information to build a unified picture of all of the dCDNs available to it. In architectures that share detailed capability information, the uCDN could perform the entire request-routing operation down to selecting a particular surrogate in the dCDN. However, when the uCDN needs to deal with many potential dCDNs, this approach does not scale, especially for dCDNs with thousands or tens of thousands of surrogates; the volume of updates to footprint and capability becomes onerous.

Were the volume of FCI updates from dCDNs to exceed the volume of requests to the uCDN, it might make more sense for the uCDN to query dCDNs upon receiving requests (as is the case in the recursive redirection mode described in [RFC7336]), instead of receiving advertisements and tracking the state of dCDNs. The advantage of querying dCDNs would be that much of the dynamic data that dCDNs cannot share with the uCDN would now be factored into the uCDN's decision. dCDNs need not replicate any state to the uCDN - uCDNs could effectively operate in a stateless mode.

The semantics of both footprint and capability advertisement depend on the service model here: are there cases where a synchronous query/response model would work better for the uCDN decision than a state replication model?

2.4. Avoiding or Handling 'cheating' dCDNs

In a situation where more than one dCDN is willing to serve a given end user request, it might be attractive for a dCDN to 'cheat' in the sense that the dCDN provides inaccurate information to the uCDN in order to convince the uCDN to select it over 'competing' dCDNs. It could therefore be desirable to take away the incentive for dCDNs to cheat (in information advertised) as much as possible. One option is to make the information the dCDN advertises somehow verifiable for

the uCDN. On the other hand, a cheating dCDN might be avoided or handled by the fact that there will be strong contractual agreements between a uCDN and a dCDN, so that a dCDN would risk severe penalties or legal consequences when caught cheating.

Overall, the information a dCDN advertises (in the long run) needs to be somehow qualitatively verifiable by the uCDN, though possibly through non-real-time out-of-band audits. It is probably an overly strict requirement to mandate that such verification be possible "immediately", i.e., during the request routing process itself. If the uCDN can detect a cheating dCDN at a later stage, it might suffice for the uCDN to "de-incentivize" cheating because it would negatively affect the long-term business relationship with a particular dCDN.

3. Focusing on Capabilities with Footprint Restrictions

Given the design considerations listed in the previous section, it seems reasonable to assume that in most cases it is the uCDN that makes the decision on selecting a certain dCDN for request routing based on information the uCDN has received from this particular dCDN. It can be assumed that 'cheating' CDNs will be dealt with via means outside the scope of CDNI and that the information advertised between CDNs is accurate. In addition, excluding the use of qualitative information (e.g., surrogate proximity, delivery latency, surrogate load) to predict the quality of delivery would further simplify the use case allowing it to better focus on the basic functionality of the FCI.

Further understanding that in most cases contractual agreements will define the basic coverage used in delegation decisions, the primary focus of FCI is on providing updates to the basic capabilities and coverage by the dCDNs. As such, FCI has chosen the semantics of "capabilities with footprint restrictions".

4. Footprint and Capabilities Extension

Other optional "coverage/reachability" types of footprint or "resource" types of footprint may be defined by future specifications. To facilitate this, a clear process for specifying optional footprint types in an IANA registry is specified in the CDNI Metadata Footprint Types registry (defined in the CDNI Metadata Interface document [I-D.ietf-cdni-metadata]).

This document also registers CDNI Payload Types [RFC7736] for the initial capability types (see Section 6):

- o Delivery Protocol (for delivering content to the end user)

- o Acquisition Protocol (for acquiring content from the uCDN or origin server)
- o Redirection Mode (e.g., DNS Redirection vs. HTTP Redirection as discussed in [RFC7336])
- o CDNI Logging (i.e., supported logging fields)
- o CDNI Metadata (i.e., supported Generic Metadata types)

Each payload type is prefaced with "FCI.". Updates to capability objects MUST indicate the version of the capability object in a newly registered payload type, e.g., by appending ".v2". Each capability type MAY have a list of valid values. Future specifications which define a given capability MUST define any necessary registries (and the rules for adding new entries to the registry) for the values advertised for a given capability type.

The "CDNI Logging record-types" registry [I-D.ietf-cdni-logging] defines all known record types, including mandatory-to-implement record-types. Advertising support for mandatory-to-implement record-types would be redundant. CDNs SHOULD NOT advertise support for mandatory-to-implement record-types.

The "CDNI Logging Fields Names" registry [I-D.ietf-cdni-logging] defines all known logging fields. Logging fields may be reused by different record-types and be mandatory-to-implement in some record-types, but optional in other record-types. CDNs MUST advertise support for optional logging fields within the context of a specific record-type. CDNs SHOULD NOT advertise support for mandatory-to-implement logging fields, for a given record-type. The following logging fields are defined as optional for the "cdni_http_request_v1" record-type in the CDNI Logging Interface document [I-D.ietf-cdni-logging]:

- o s-ccid
- o s-sid

The CDNI Metadata Interface document [I-D.ietf-cdni-metadata] requires that CDNs be able to parse all the defined metadata objects, but does not require dCDNs to support enforcement of non-structural GenericMetadata objects. Advertising support for mandatory-to-enforce GenericMetadata types MUST be supported. Advertising support for non-mandatory-to-enforce GenericMetadata types SHOULD be supported. Advertisement of non-mandatory-to-enforce GenericMetadata MAY be necessary, e.g., to signal temporary outages and subsequent

recovery. It is expected that structural metadata will be supported at all times.

The notion of optional types of footprint and capabilities implies that certain implementations might not support all kinds of footprint and capabilities. Therefore, any FCI solution protocol MUST define how the support for optional types of footprint/capabilities will be negotiated between a uCDN and a dCDN that use the particular FCI protocol. In particular, any FCI solution protocol MUST specify how to handle failure cases or non-supported types of footprint/capabilities.

In general, a uCDN MAY ignore capabilities or types of footprints it does not understand; in this case it only selects a suitable dCDN based on the types of capabilities and footprint it understands. Similarly, if a dCDN does not use an optional capability or footprint which is, however, supported by a uCDN, this causes no problem for the FCI functionality because the uCDN decides on the remaining capabilities/footprint information that is being conveyed by the dCDN.

5. Capability Advertisement Object

To support extensibility, the FCI defines a generic base object (similar to the CDNI Metadata interface GenericMetadata object) [I-D.ietf-cdni-metadata] to facilitate a uniform set of mandatory parsing requirements for all future FCI objects.

Future object definitions (e.g. regarding CDNI Metadata or Logging) will build off the base object defined here, but will be specified in separate documents.

Note: In the following sections, the term "mandatory-to-specify" is used to convey which properties MUST be included when serializing a given capability object. When mandatory-to-specify is defined as "Yes" for an individual property, it means that if the object containing that property is included in an FCI message, then the mandatory-to-specify property MUST also be included.

5.1. Base Advertisement Object

The FCIBase object is an abstraction for managing individual CDNI capabilities in an opaque manner.

Property: capability-type

Description: CDNI Capability object type.

Type: FCI specific CDNI Payload type (from the CDNI Payload Types registry [RFC7736])

Mandatory-to-Specify: Yes.

Property: capability-value

Description: CDNI Capability object.

Type: Format/Type is defined by the value of capability-type property above.

Mandatory-to-Specify: Yes.

Property: footprints

Description: CDNI Capability Footprint.

Type: List of CDNI Footprint objects (as defined in [I-D.ietf-cdni-metadata]).

Mandatory-to-Specify: No.

5.2. Encoding

CDNI FCI objects MUST be encoded using JSON [RFC7159] and MUST also follow the recommendations of I-JSON [RFC7493]. FCI objects are composed of a dictionary of (key,value) pairs where the keys are the property names and the values are the associated property values.

The keys of the dictionary are the names of the properties associated with the object and are therefore dependent on the specific object being encoded (i.e., dependent on the CDNI Payload Type of the capability or the CDNI Metadata Footprint Type of the footprint). Likewise, the values associated with each property (dictionary key) are dependent on the specific object being encoded (i.e., dependent on the CDNI Payload Type of the capability or the CDNI Metadata Footprint Type of the footprint).

Dictionary keys (properties) in JSON are case sensitive. By convention, any dictionary key (property) defined by this document MUST be lowercase.

5.3. Delivery Protocol Capability Object

The Delivery Protocol capability object is used to indicate support for one or more of the protocols listed in the CDNI Metadata Protocol

Types registry (defined in the CDNI Metadata Interface document [I-D.ietf-cdni-metadata]).

Property: delivery-protocols

Description: List of supported CDNI Delivery Protocols.

Type: List of Protocol Types (from the CDNI Metadata Protocol Types registry [I-D.ietf-cdni-metadata])

Mandatory-to-Specify: Yes.

5.3.1. Delivery Protocol Capability Object Serialization

The following shows an example of Delivery Protocol Capability Object Serialization, for a CDN that supports only HTTP/1.1 without TLS for content delivery.

```
{
  "capabilities": [
    {
      "capability-type": "FCI.DeliveryProtocol",
      "capability-value": {
        "delivery-protocols": [
          "http/1.1",
        ]
      },
      "footprints": [
        <Footprint objects>
      ]
    }
  ]
}
```

5.4. Acquisition Protocol Capability Object

The Acquisition Protocol capability object is used to indicate support for one or more of the protocols listed in the CDNI Metadata Protocol Types registry (defined in the CDNI Metadata Interface document [I-D.ietf-cdni-metadata]).

Property: acquisition-protocols

Description: List of supported CDNI Acquisition Protocols.

Type: List of Protocol Types (from the CDNI Metadata Protocol Types registry [I-D.ietf-cdni-metadata])

Mandatory-to-Specify: Yes.

5.4.1. Acquisition Protocol Capability Object Serialization

The following shows an example of Acquisition Protocol Capability Object Serialization, for a CDN that supports HTTP/1.1 with or without TLS for content acquisition.

```
{
  "capabilities": [
    {
      "capability-type": "FCI.AcquisitionProtocol",
      "capability-value": {
        "acquisition-protocols": [
          "http/1.1",
          "https/1.1"
        ]
      },
      "footprints": [
        <Footprint objects>
      ]
    }
  ]
}
```

5.5. Redirection Mode Capability Object

The Redirection Mode capability object is used to indicate support for one or more of the modes listed in the CDNI Capabilities Redirection Modes registry (see Section 6.2).

Property: redirection-modes

Description: List of supported CDNI Redirection Modes.

Type: List of Redirection Modes (from Section 6.2)

Mandatory-to-Specify: Yes.

5.5.1. Redirection Mode Capability Object Serialization

The following shows an example of Redirection Mode Capability Object Serialization, for a CDN that supports only iterative (but not recursive) redirection with HTTP and DNS.

```

{
  "capabilities": [
    {
      "capability-type": "FCI.RedirectionMode",
      "capability-value": {
        "redirection-modes": [
          "DNS-I",
          "HTTP-I"
        ]
      }
    },
    {
      "footprints": [
        <Footprint objects>
      ]
    }
  ]
}

```

5.6. CDNI Logging Capability Object

The CDNI Logging capability object is used to indicate support for CDNI Logging record-types, as well as CDNI Logging fields which are marked as optional for the specified record-types [I-D.ietf-cdni-logging].

Property: record-type

Description: Supported CDNI Logging record-type.

Type: String corresponding to an entry from the CDNI Logging record-types registry [I-D.ietf-cdni-logging])

Mandatory-to-Specify: Yes.

Property: fields

Description: List of supported CDNI Logging fields that are optional for the specified record-type.

Type: List of Strings corresponding to entries from the CDNI Logging Field Names registry [I-D.ietf-cdni-logging].

Mandatory-to-Specify: No. Default is that all optional fields are supported. Omission of this field MUST be interpreted as "all optional fields are supported". An empty list MUST be interpreted as "no optional fields are supported. Otherwise, if a list of fields is provided, the fields in that list MUST be interpreted as "the only optional fields that are supported".

5.6.1. CDNI Logging Capability Object Serialization

The following shows an example of CDNI Logging Capability Object Serialization, for a CDN that supports the optional Content Collection ID logging field (but not the optional Session ID logging field) for the "cdni_http_request_v1" record type.

```
{
  "capabilities": [
    {
      "capability-type": "FCI.Logging",
      "capability-value": {
        "record-type": "cdni_http_request_v1",
        "fields": [ "s-ccid" ]
      },
      "footprints": [
        <Footprint objects>
      ]
    }
  ]
}
```

The next example shows the CDNI Logging Capability Object Serialization, for a CDN that supports all optional fields for the "cdni_http_request_v1" record type.

```
{
  "capabilities": [
    {
      "capability-type": "FCI.Logging",
      "capability-value": {
        "record-type": "cdni_http_request_v1"
      },
      "footprints": [
        <Footprint objects>
      ]
    }
  ]
}
```

5.7. CDNI Metadata Capability Object

The CDNI Metadata capability object is used to indicate support for CDNI GenericMetadata types [I-D.ietf-cdni-metadata].

Property: metadata

Description: List of supported CDNI GenericMetadata types.

Type: List of Strings corresponding to entries from the CDNI Payload Type registry [RFC7736]) that correspond to CDNI GenericMetadata objects.

Mandatory-to-Specify: Yes. An empty list MUST be interpreted as "no GenericMetadata types are supported", i.e., "only structural metadata and simple types are supported"; otherwise, the list must be interpreted as containing "the only GenericMetadata types that are supported" (in addition to structural metadata and simple types) [I-D.ietf-cdni-metadata].

5.7.1. CDNI Metadata Capability Object Serialization

The following shows an example of CDNI Metadata Capability Object Serialization, for a CDN that supports only the SourceMetadata GenericMetadata type (i.e., it can acquire and deliver content, but cannot enforce and security policies, e.g., time, location, or protocol ACLs).

```
{
  "capabilities": [
    {
      "capability-type": "FCI.Metadata",
      "capability-value": {
        "metadata": ["MI.SourceMetadata"]
      },
      "footprints": [
        <Footprint objects>
      ]
    }
  ]
}
```

The next example shows the CDNI Metadata Capability Object Serialization, for a CDN that supports only structural metadata (i.e., it can parse metadata as a transit CDN, but cannot enforce security policies or deliver content).

```

{
  "capabilities": [
    {
      "capability-type": "FCI.Metadata",
      "capability-value": {
        "metadata": []
      },
      "footprints": [
        <Footprint objects>
      ]
    }
  ]
}

```

6. IANA Considerations

6.1. CDNI Payload Types

This document requests the registration of the following CDNI Payload Types under the IANA CDNI Payload Type registry:

Payload Type	Specification
FCI.DeliveryProtocol	RFCthis
FCI.AcquisitionProtocol	RFCthis
FCI.RedirectionMode	RFCthis
FCI.Logging	RFCthis
FCI.Metadata	RFCthis

[RFC Editor: Please replace RFCthis with the published RFC number for this document.]

6.1.1. CDNI FCI DeliveryProtocol Payload Type

Purpose: The purpose of this payload type is to distinguish FCI advertisement objects for supported delivery protocols

Interface: FCI

Encoding: see Section 5.3

6.1.2. CDNI FCI AcquisitionProtocol Payload Type

Purpose: The purpose of this payload type is to distinguish FCI advertisement objects for supported acquisition protocols

Interface: FCI

Encoding: see Section 5.4

6.1.3. CDNI FCI RedirectionMode Payload Type

Purpose: The purpose of this payload type is to distinguish FCI advertisement objects for supported redirection modes

Interface: FCI

Encoding: see Section 5.5

6.1.4. CDNI FCI Logging Payload Type

Purpose: The purpose of this payload type is to distinguish FCI advertisement objects for supported CDNI Logging record-types and optional CDNI Logging Field Names.

Interface: FCI

Encoding: see Section 5.6

6.1.5. CDNI FCI Metadata Payload Type

Purpose: The purpose of this payload type is to distinguish FCI advertisement objects for supported CDNI GenericMetadata types.

Interface: FCI

Encoding: see Section 5.7

6.2. Redirection Mode Registry

The IANA is requested to create a new "CDNI Capabilities Redirection Modes" registry in the "Content Delivery Networks Interconnection (CDNI) Parameters" category. The "CDNI Capabilities Redirection Modes" namespace defines the valid redirection modes that can be advertised as supported by a CDN. Additions to the Redirection Mode namespace conform to the "IETF Review" policy as defined in [RFC5226].

The following table defines the initial Redirection Modes:

Redirection Mode	Description	RFC
DNS-I	Iterative DNS-based Redirection	RFCthis
DNS-R	Recursive DNS-based Redirection	RFCthis
HTTP-I	Iterative HTTP-based Redirection	RFCthis
HTTP-R	Recursive HTTP-based Redirection	RFCthis

[RFC Editor: Please replace RFCthis with the published RFC number for this document.]

7. Security Considerations

This specification describes the semantics for capabilities and footprint advertisement objects across interconnected CDNs. It does not, however, specify a concrete protocol for transporting those objects. Specific security mechanisms can only be selected for concrete protocols that instantiate these semantics. This document does, however, place some high-level security constraints on such protocols.

All protocols that implement these semantics are REQUIRED to provide integrity and authentication services. Without authentication and integrity, an attacker could trivially deny service by forging a footprint advertisement from a dCDN which claims the network has no footprint or capability. This would prevent the uCDN from delegating any requests to the dCDN. Since a pre-existing relationship between all dCDNs and uCDNs is assumed by CDNI, the exchange of any necessary credentials could be conducted before the FCI interface is brought online. The authorization decision to accept advertisements would also follow this pre-existing relationship and any contractual obligations that it stipulates.

All protocols that implement these semantics are REQUIRED to provide confidentiality services. Some dCDNs are willing to share information about their footprint or capabilities with a uCDN but not with other, competing dCDNs. For example, if a dCDN incurs an outage that reduces footprint coverage temporarily, that could be information the dCDN would want to share confidentially with the uCDN.

As specified in this document, the security requirements of the FCI could be met by transport-layer security mechanisms coupled with domain certificates as credentials (e.g., TLS transport for HTTP as

per [RFC2818] and [RFC7230], with usage guidance from [RFC7525]) between CDNs. There is no apparent need for further object-level security in this framework, as the trust relationships it defines are bilateral relationships between uCDNs and dCDNs rather than transitive relationships.

8. References

8.1. Normative References

- [I-D.ietf-cdni-logging]
Faucheur, F., Bertrand, G., Oprescu, I., and R. Peterkofsky, "CDNI Logging Interface", draft-ietf-cdni-logging-25 (work in progress), April 2016.
- [I-D.ietf-cdni-metadata]
Niven-Jenkins, B., Murray, R., Caulfield, M., and K. Ma, "CDN Interconnection Metadata", draft-ietf-cdni-metadata-16 (work in progress), April 2016.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, DOI 10.17487/RFC5226, May 2008, <<http://www.rfc-editor.org/info/rfc5226>>.
- [RFC7159] Bray, T., Ed., "The JavaScript Object Notation (JSON) Data Interchange Format", RFC 7159, DOI 10.17487/RFC7159, March 2014, <<http://www.rfc-editor.org/info/rfc7159>>.
- [RFC7336] Peterson, L., Davie, B., and R. van Brandenburg, Ed., "Framework for Content Distribution Network Interconnection (CDNI)", RFC 7336, DOI 10.17487/RFC7336, August 2014, <<http://www.rfc-editor.org/info/rfc7336>>.
- [RFC7493] Bray, T., Ed., "The I-JSON Message Format", RFC 7493, DOI 10.17487/RFC7493, March 2015, <<http://www.rfc-editor.org/info/rfc7493>>.

8.2. Informative References

- [RFC2818] Rescorla, E., "HTTP Over TLS", RFC 2818, DOI 10.17487/RFC2818, May 2000, <<http://www.rfc-editor.org/info/rfc2818>>.

- [RFC6707] Niven-Jenkins, B., Le Faucheur, F., and N. Bitar, "Content Distribution Network Interconnection (CDNI) Problem Statement", RFC 6707, DOI 10.17487/RFC6707, September 2012, <<http://www.rfc-editor.org/info/rfc6707>>.
- [RFC6770] Bertrand, G., Ed., Stephan, E., Burbridge, T., Eardley, P., Ma, K., and G. Watson, "Use Cases for Content Delivery Network Interconnection", RFC 6770, DOI 10.17487/RFC6770, November 2012, <<http://www.rfc-editor.org/info/rfc6770>>.
- [RFC7230] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", RFC 7230, DOI 10.17487/RFC7230, June 2014, <<http://www.rfc-editor.org/info/rfc7230>>.
- [RFC7337] Leung, K., Ed. and Y. Lee, Ed., "Content Distribution Network Interconnection (CDNI) Requirements", RFC 7337, DOI 10.17487/RFC7337, August 2014, <<http://www.rfc-editor.org/info/rfc7337>>.
- [RFC7525] Sheffer, Y., Holz, R., and P. Saint-Andre, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", BCP 195, RFC 7525, DOI 10.17487/RFC7525, May 2015, <<http://www.rfc-editor.org/info/rfc7525>>.
- [RFC7736] Ma, K., "Content Delivery Network Interconnection (CDNI) Media Type Registration", RFC 7736, DOI 10.17487/RFC7736, December 2015, <<http://www.rfc-editor.org/info/rfc7736>>.

Appendix A. Main Use Case to Consider

Focusing on a main use case that contains a simple (yet somewhat challenging), realistic, and generally imaginable scenario can help in narrowing down the requirements for the CDNI FCI. To this end, the following (simplified) use case can help in clarifying the semantics of footprint and capabilities for CDNI. In particular, the intention of the use case is to clarify what information needs to be exchanged on the CDNI FCI, what types of information need to be supported in a mandatory fashion (and which can be considered optional), and what types of information need to be updated with respect to a priori established CDNI contracts.

Use case: A given uCDN has several dCDNs. It selects one dCDN for delivery protocol A and footprint 1 and another dCDN for delivery protocol B and footprint 1. The dCDN that serves delivery protocol B has a further, transitive (level-2) dCDN, that serves delivery protocol B in a subset of footprint 1 where the first-level dCDN

cannot serve delivery protocol B itself. What happens if capabilities change in the transitive level-2 dCDN that might affect how the uCDN selects a level-1 dCDN (e.g., in case the level-2 dCDN cannot serve delivery protocol B anymore)? How will these changes be conveyed to the uCDN? In particular, what information does the uCDN need to be able to select a new first-level dCDN, either for all of footprint 1 or only for the subset of footprint 1 that the transitive level-2 dCDN served on behalf of the first-level dCDN?

Appendix B. Semantics for Footprint Advertisement

Roughly speaking, "footprint" can be defined as "ability and willingness to serve" by a dCDN. However, in addition to simple "ability and willingness to serve", the uCDN could want additional information to make a dCDN selection decision, e.g., "how well" a given dCDN can actually serve a given end user request. The "ability and willingness" to serve SHOULD be distinguished from the subjective qualitative measurement of "how well" it was served. One can imagine that such additional information is implicitly associated with a given footprint, due to contractual agreements, SLAs, business relationships, or past perceptions of dCDN quality. As an alternative, such additional information could also be explicitly tagged along with the footprint.

It is reasonable to assume that a significant part of the actual footprint advertisement will happen in contractual agreements between participating CDNs, prior to the advertisement phase using the CDNI FCI. The reason for this assumption is that any contractual agreement is likely to contain specifics about the dCDN coverage (footprint) to which the contractual agreement applies. In particular, additional information to judge the delivery quality associated with a given dCDN footprint might be defined in contractual agreements, outside of the CDNI FCI. Further, one can assume that dCDN contractual agreements about the delivery quality associated with a given footprint will probably be based on high-level aggregated statistics and not too detailed.

Given that a large part of footprint advertisement will actually happen in contractual agreements, the semantics of CDNI footprint advertisement refer to answering the following question: what exactly still needs to be advertised by the CDNI FCI? For instance, updates about temporal failures of part of a footprint can be useful information to convey via the CDNI request routing interface. Such information would provide updates on information previously agreed in contracts between the participating CDNs. In other words, the CDNI FCI is a means for a dCDN to provide changes/updates regarding a footprint it has prior agreed to serve in a contract with a uCDN.

Generally speaking, one can imagine two categories of footprint to be advertised by a dCDN:

- o Footprint could be defined based on "coverage/reachability", where coverage/reachability refers to a set of prefixes, a geographic region, or similar boundary. The dCDN claims that it can cover/reach 'end user requests coming from this footprint'.
- o Footprint could be defined based on "resources", where resources refers to surrogates a dCDN claims to have (e.g., the location of surrogates/resources). The dCDN claims that 'from this footprint' it can serve incoming end user requests.

For each of these footprint types, there are capabilities associated with a given footprint:

- o capabilities such as delivery protocol, redirection mode, and metadata, which are supported in the coverage area for a "coverage/reachability" defined footprint, or
- o capabilities of resources, such as delivery protocol, redirection mode, and metadata, which apply to a "resource" defined footprint.

"Resource" types of footprints are more specific than "coverage/reachability" types of footprints, where the actual coverage/reachability are extrapolated from the resource location (e.g., netmask applied to resource IP address to derive IP-prefix). The specific methods for extrapolating coverage/reachability from resource location are beyond the scope of this document. In the degenerate case, the resource address could be specified as a coverage/reachability type of footprint, in which case no extrapolation is necessary. Resource types of footprints could expose the internal structure of a CDN network which could be undesirable. As such, the resource types of footprints are not considered mandatory to support for CDNI.

Footprints can be viewed as constraints for delegating requests to a dCDN: A dCDN footprint advertisement tells the uCDN the limitations for delegating a request to the dCDN. For IP prefixes or ASN(s), the footprint signals to the uCDN that it should consider the dCDN a candidate only if the IP address of the request routing source falls within the prefix set (or ASN, respectively). The CDNI specifications do not define how a given uCDN determines what address ranges are in a particular ASN. Similarly, for country codes a uCDN should only consider the dCDN a candidate if it covers the country of the request routing source. The CDNI specifications do not define how a given uCDN determines the country of the request routing source. Multiple footprint constraints are additive: the

advertisement of different types of footprint narrows the dCDN candidacy cumulatively.

Independent of the exact type of a footprint, a footprint might also include the connectivity of a given dCDN to other CDNs that are able to serve content to users on behalf of that dCDN, to cover cases with cascaded CDNs. Further, the dCDN needs to be able to express its footprint to an interested uCDN in a comprehensive form, e.g., as a data set containing the complete footprint. Making incremental updates, however, to express dynamic changes in state is also desirable.

Appendix C. Semantics for Capabilities Advertisement

In general, the dCDN needs to be able to express its general capabilities to the uCDN. These general capabilities could express if the dCDN supports a given service, for instance, HTTP vs HTTPS delivery. Furthermore, the dCDN needs to be able to express particular capabilities for the delivery in a particular footprint area. For example, the dCDN might in general offer HTTPS but not in some specific areas, either for maintenance reasons or because the surrogates covering this particular area cannot deliver this type of service. Hence, in certain cases footprint and capabilities are tied together and cannot be interpreted independently from each other. In such cases, i.e., where capabilities need to be expressed on a per footprint basis, it could be beneficial to combine footprint and capabilities advertisement.

A high-level and very rough semantic for capabilities is thus the following: Capabilities are types of information that allow a uCDN to determine if a dCDN is able (and willing) to accept (and properly handle) a delegated content request. In addition, Capabilities are characterized by the fact that this information can change over time based on the state of the network or surrogates.

At a first glance, several broad categories of capabilities seem useful to convey via an advertisement interface, however, advertising capabilities that change highly dynamically (e.g., real-time delivery performance metrics, CDN resource load, or other highly dynamically changing QoS information) is beyond the scope for CDNI FCI. First, out of the multitude of possible metrics and capabilities, it is hard to agree on a subset and the precise metrics to be used. Second, it seems infeasible to specify such highly dynamically changing capabilities and the corresponding metrics within a reasonable time-frame.

Useful capabilities refer to information that does not change highly dynamically and which in many cases is absolutely necessary to decide

on a particular dCDN for a given end user request. For instance, if an end user request concerns the delivery of a video file with a certain protocol, the uCDN needs to know if a given dCDN has the capability of supporting this delivery protocol.

Similar to footprint advertisement, it is reasonable to assume that a significant part of the actual (resource) capabilities advertisement will happen in contractual agreements between participating CDNs, i.e., prior to the advertisement phase using the CDNI FCI. The role of capability advertisement is hence rather to enable the dCDN to update a uCDN on changes since a contract has been set up (e.g., in case a new delivery protocol is suddenly being added to the list of supported delivery protocols of a given dCDN, or in case a certain delivery protocol is suddenly not being supported anymore due to failures). Capabilities advertisement thus refers to conveying information to a uCDN about changes/updates of certain capabilities with respect to a given contract.

Given these semantics, it needs to be decided what exact capabilities are useful and how these can be expressed. Since the details of CDNI contracts are not known at the time of this writing (and the CDNI interface are better off being agnostic to these contracts anyway), it remains to be seen what capabilities will be used to define agreements between CDNs in practice. One implication for standardization could be to initially only specify a very limited set of mandatory capabilities for advertisement and have on top of that a flexible data model that allows exchanging additional capabilities when needed. Still, agreement needs to be found on which capabilities (if any) will be mandatory among CDNs.

It is not feasible to enumerate all the possible options for the mandatory capabilities listed above (e.g., all the potential delivery protocols or metadata options) or anticipate all the future needs for additional capabilities. It would be unreasonable to burden the CDNI FCI specification with defining each supported capability. Instead, the CDNI FCI specification should define a generic protocol for conveying any capability information (e.g. with common encoding, error handling, and security mechanism; further requirements for the CDNI FCI Advertisement Interface are listed in [RFC7337]).

Appendix D. Acknowledgment

Jan Seedorf is partially supported by the GreenICN project (GreenICN: Architecture and Applications of Green Information Centric Networking), a research project supported jointly by the European Commission under its 7th Framework Program (contract no. 608518) and the National Institute of Information and Communications Technology (NICT) in Japan (contract no. 167). The views and conclusions

contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the GreenICN project, the European Commission, or NICT.

Martin Stiernerling provided initial input to this document and valuable comments to the ongoing discussions among the authors of this document. Thanks to Francois Le Faucheur and Scott Wainner for providing valuable comments and suggestions to the text.

Authors' Addresses

Jan Seedorf
NEC
Kurfuerstenanlage 36
Heidelberg 69115
Germany

Phone: +49 6221 4342 221
Fax: +49 6221 4342 155
Email: seedorf@neclab.eu

Jon Peterson
NeuStar
1800 Sutter St Suite 570
Concord CA 94520
USA

Email: jon.peterson@neustar.biz

Stefano Previdi
Cisco Systems
Via Del Serafico 200
Rome 0144
Italy

Email: sprevidi@cisco.com

Ray van Brandenburg
TNO
Brassersplein 2
Delft 2612CT
The Netherlands

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

Kevin J. Ma
Ericsson
43 Nagog Park
Acton, MA 01720
USA

Phone: +1 978-844-5100
Email: kevin.j.ma@ericsson.com

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: April 21, 2014

F. Le Faucheur, Ed.
Cisco Systems
G. Bertrand, Ed.
I. Oprescu, Ed.
Orange
R. Peterkofsky
Skytide, Inc.
October 18, 2013

CDNI Logging Interface
draft-ietf-cdni-logging-08

Abstract

This memo specifies the Logging interface between a downstream CDN (dCDN) and an upstream CDN (uCDN) that are interconnected as per the CDN Interconnection (CDNI) framework. First, it describes a reference model for CDNI logging. Then, it specifies the CDNI Logging File format and the actual protocol for exchange of CDNI Logging Files.

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 April 21, 2014.

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	3
1.1. Terminology	3
1.2. Requirements Language	4
2. CDNI Logging Reference Model	5
2.1. CDNI Logging interactions	5
2.2. Overall Logging Chain	8
2.2.1. Logging Generation and During-Generation Aggregation	9
2.2.2. Logging Collection	10
2.2.3. Logging Filtering	10
2.2.4. Logging Rectification and Post-Generation Aggregation	11
2.2.5. Log-Consuming Applications	11
2.2.5.1. Maintenance/Debugging	11
2.2.5.2. Accounting	12
2.2.5.3. Analytics and Reporting	12
2.2.5.4. Security	12
2.2.5.5. Legal Logging Duties	12
2.2.5.6. Notions common to multiple Log Consuming	
Applications	13
3. CDNI Logging File	15
3.1. Rules	15
3.2. CDNI Logging File Structure	16
3.3. CDNI Logging File Directives	18
3.4. CDNI Logging Records	21
3.4.1. HTTP Request Logging Record	22
3.5. CDNI Logging File Example	29
4. CDNI Logging File Exchange Protocol	30
4.1. CDNI Logging Feed	30
4.1.1. Atom Formatting	31
4.1.2. Updates to Log Files and the Feed	31
4.1.3. Redundant Feeds	32
4.1.4. Example CDNI Logging Feed	32
4.2. CDNI Logging File Pull	33
5. IANA Considerations	35
5.1. CDNI Logging Directive Names Registry	35
5.2. CDNI Logging Record-Types Registry	35
5.3. CDNI Logging Field Names Registry	36
5.4. CDNI Logging MIME Media Type	37
6. Security Considerations	37
6.1. Authentication, Confidentiality, Integrity Protection . .	37
6.2. Denial of Service	38

6.3. Privacy	38
7. Acknowledgments	38
8. References	39
8.1. Normative References	39
8.2. Informative References	39
Appendix A. Compliance with CDNI Requirements	40
Authors' Addresses	45

1. Introduction

This memo specifies the CDNI Logging interface between a downstream CDN (dCDN) and an upstream CDN (uCDN). First, it describes a reference model for CDNI logging. Then, it specifies the CDNI Logging File format and the actual protocol for exchange of CDNI Logging Files.

The reader should be familiar with the following documents:

- o CDNI problem statement [RFC6707] and framework [I-D.ietf-cdni-framework] identify a Logging interface,
- o Section 8 of [I-D.ietf-cdni-requirements] specifies a set of requirements for Logging,
- o [RFC6770] outlines real world use-cases for interconnecting CDNs. These use cases require the exchange of Logging information between the dCDN and the uCDN.

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

The present document describes:

- o The CDNI Logging reference model (Section 2),
- o The CDNI Logging File format (Section 3),
- o The CDNI Logging File Exchange protocol (Section 4).

1.1. Terminology

In this document, the first letter of each CDNI-specific term is capitalized. We adopt the terminology described in [RFC6707] and [I-D.ietf-cdni-framework], and extend it with the additional terms defined below.

For clarity, we use the word "Log" only for referring to internal CDN logs and we use the word "Logging" for any inter-CDN information

exchange and processing operations related to CDNI Logging interface. Log and Logging formats may be different.

CDN Logging information: logging information generated and collected within a CDN

CDNI Logging information: logging information exchanged across CDNs using the CDNI Logging Interface

Logging information: logging information generated and collected within a CDN or obtained from another CDN using the CDNI Logging Interface

CDNI Logging Field: an atomic element of information that can be included in a CDNI Logging Record. The time an event/task started, the IP address of an End user to whom content was delivered, and the URI of the content delivered are examples of CDNI Logging Fields.

CDNI Logging Record: an information record providing information about a specific event. This comprises a collection of CDNI Logging Fields.

CDNI Logging File: a file containing CDNI Logging Records, as well as additional information facilitating the processing of the CDNI Logging Records.

CDN Reporting: the process of providing the relevant information that will be used to create a formatted content delivery report provided to the CSP in deferred time. Such information typically includes aggregated data that can cover a large period of time (e.g., from hours to several months). Uses of Reporting include the collection of charging data related to CDN services and the computation of Key Performance Indicators (KPIs).

CDN Monitoring: the process of providing content delivery information in real-time. Monitoring typically includes data in real time to provide visibility of the deliveries in progress, for service operation purposes. It presents a view of the global health of the services as well as information on usage and performance, for network services supervision and operation management. In particular, monitoring data can be used to generate alarms.

1.2. 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].

2. CDNI Logging Reference Model

2.1. CDNI Logging interactions

The CDNI logging reference model between a given uCDN and a given dCDN involves the following interactions:

- o customization by the uCDN of the CDNI logging information to be provided by the dCDN to the uCDN (e.g. control of which logging fields are to be communicated to the uCDN for a given task performed by the dCDN, control of which types of events are to be logged). The dCDN takes into account this CDNI logging customization information to determine what logging information to provide to the uCDN, but it may, or may not, take into account this CDNI logging customization information to influence what CDNI logging information is to be generated and collected within the dCDN (e.g. even if the uCDN requests a restricted subset of the logging information, the dCDN may elect to generate a broader set of logging information). The mechanism to support the customisation by the uCDN of CDNI Logging information is outside the scope of this document and left for further study. We note that the CDNI Control interface or the CDNI Metadata interface appear as candidate interfaces on which to potentially build such a customisation mechanism in the future. Before such a mechanism is available, the uCDN and dCDN are expected to agree off-line on what CDNI logging information is to be provide by dCDN to UCDN and rely on management plane actions to configure the CDNI Logging functions to generate (respectively, expect) in dCDN (respectively, in uCDN).
- o generation and collection by the dCDN of logging information related to the completion of any task performed by the dCDN on behalf of the uCDN (e.g., delivery of the content to an end user) or related to events happening in the dCDN that are relevant to the uCDN (e.g., failures or unavailability in dCDN). This takes place within the dCDN and does not directly involve CDNI interfaces.
- o communication by the dCDN to the uCDN of the logging information collected by the dCDN relevant to the uCDN. This is supported by the CDNI Logging interface and in the scope of the present document. For example, the uCDN may use this logging information to charge the CSP, to perform analytics and monitoring for operational reasons, to provide analytics and monitoring views on its content delivery to the CSP or to perform trouble-shooting. This document exclusively specifies non-real-time exchange of logging information. Closer to real-time exchange of logging information (say sub-minute or sub-second) is outside the scope of


the present document and left for further study. This document exclusively specifies exchange of logging information related to content delivery. Exchange of logging information related to operational events (e.g. dCDN request routing function unavailable, content acquisition failure by dCDN) for audit or operational reactive adjustments by uCDN is outside the scope of the present document and left for further study.

- o customization by the dCDN of the logging to be performed by the uCDN on behalf of the dCDN. The mechanism to support the customisation by the dCDN of CDNI Logging information is outside the scope of this document and left for further study.
- o generation and collection by the uCDN of logging information related to the completion of any task performed by the uCDN on behalf of the dCDN (e.g., serving of content by uCDN to dCDN for acquisition purposes by dCDN) or related to events happening in the uCDN that are relevant to the dCDN. This takes place within the uCDN and does not directly involve CDNI interfaces.
- o communication by the uCDN to the dCDN of the logging information collected by the uCDN relevant to the dCDN. For example, the dCDN might potentially benefit from this information for security auditing or content acquisition troubleshooting. This is outside the scope of this document and left for further study.

Figure 1 provides an example of CDNI Logging interactions (focusing only on the interactions that are in the scope of this document) in a particular scenario where 4 CDNs are involved in the delivery of content from a given CSP: the uCDN has a CDNI interconnection with dCDN-1 and dCDN-2. In turn, dCDN2 has a CDNI interconnection with dCDN3. In this example, uCDN, dCDN-1, dCDN-2 and dCDN-3 all participate in the delivery of content for the CSP. In this example, the CDNI Logging interface enables the uCDN to obtain logging information from all the dCDNs involved in the delivery. In the example, uCDN uses the Logging data:

- o to analyze the performance of the delivery operated by the dCDNs and to adjust its operations after the fact (e.g., request routing) as appropriate,
- o to provide (non real-time) reporting and monitoring information to CSP.

For instance, uCDN merges Logging data, extracts relevant KPIs, and presents a formatted report to the CSP, in addition to a bill for the content delivered by uCDN itself or by its dCDNs on his behalf. uCDN may also provide Logging data as raw log files to the CSP, so that the CSP can use its own logging analysis tools.



==> CDNI Logging Interface
 ***> outside the scope of CDNI

Figure 1: Interactions in CDNI Logging Reference Model

A dCDN (e.g., dCDN-2) integrates the relevant logging information obtained from its dCDNs (e.g., dCDN-3) in the logging information that it provides to the uCDN, so that the uCDN ultimately obtains all logging information relevant to a CSP for which it acts as the authoritative CDN.

Note that the format of Logging information that a CDN provides over the CDNI interface might be different from the one that the CDN uses internally. In this case, the CDN needs to reformat the Logging information before it provides this information to the other CDN over the CDNI Logging interface. Similarly, a CDN might reformat the Logging data that it receives over the CDNI Logging interface before injecting it into its log-consuming applications or before providing some of this logging information to the CSP. Such reformatting operations introduce latency in the logging distribution chain and

introduce a processing burden. Therefore, there are benefits in specifying CDNI Logging format that are suitable for use inside CDNs and also are close to the CDN Log formats commonly used in CDNs today.

2.2. Overall Logging Chain

This section discusses the overall logging chain within and across CDNs to clarify how CDN Logging information is expected to fit in this overall chain. Figure 2 illustrates the overall logging chain within the dCDN, across CDNs using the CDNI Logging interface and within the uCDN. Note that the logging chain illustrated in the Figure is obviously only an example and varies depending on the specific environments. For example, there may be more or less instantiations of each entity (i.e., there may be 4 Log consuming applications in a given CDN). As another example, there may be one instance of Rectification process per Log Consuming Application instead of a shared one.

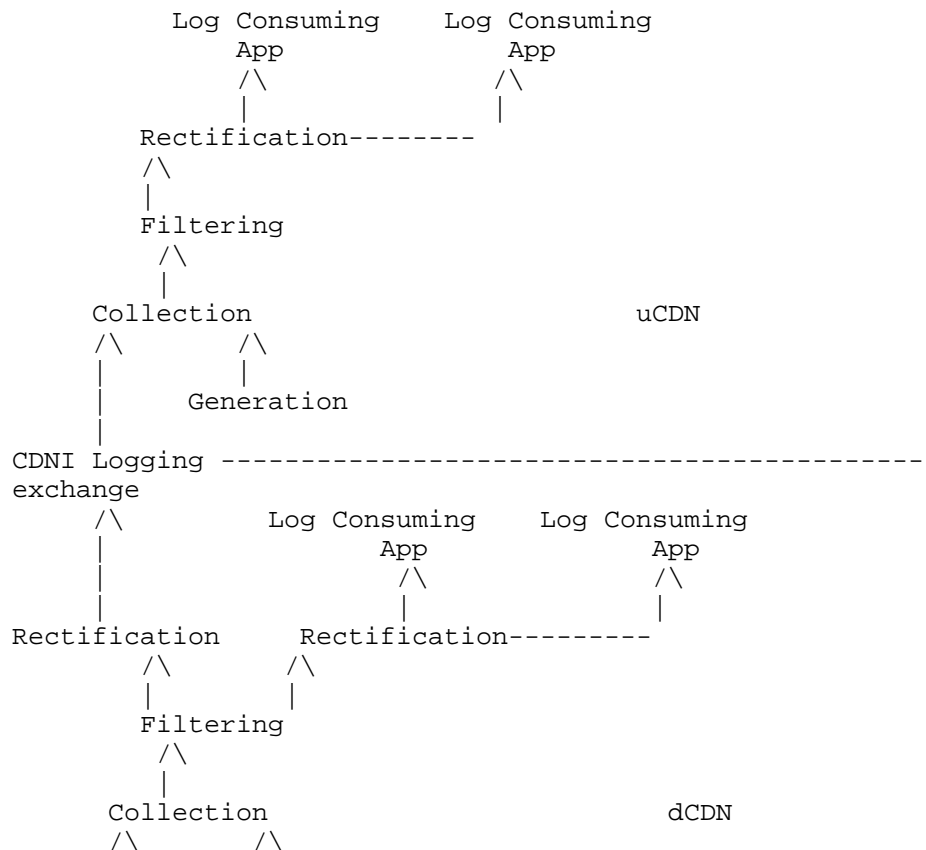




Figure 2: CDNI Logging in the overall Logging Chain

The following subsections describe each of the processes potentially involved in the logging chain of Figure 2.

2.2.1. Logging Generation and During-Generation Aggregation

CDNs typically generate logging information for all significant task completions, events, and failures. Logs are typically generated by many devices in the CDN including the surrogates, the request routing system, and the control system.

The amount of Logging information generated can be huge. Therefore, during contract negotiations, interconnected CDNs often agree on a Logging retention duration, and optionally, on a maximum size of the Logging data that the dCDN must keep. If this size is exceeded, the dCDN must alert the uCDN but may not keep more Logs for the considered time period. In addition, CDNs may aggregate logs and transmit only summaries for some categories of operations instead of the full Logging data. Note that such aggregation leads to an information loss, which may be problematic for some usages of Logging (e.g., debugging).

[RFC6983] discusses logging for HTTP Adaptive Streaming (HAS). In accordance with the recommendations articulated there, it is expected that a surrogate will generate separate logging information for delivery of each chunk of HAS content. This ensures that separate logging information can then be provided to interconnected CDNs over the CDNI Logging interface. Still in line with the recommendations of [RFC6983], the logging information for per-chunk delivery may include some information (a Content Collection IDentifier and a Session IDentifier) intended to facilitate subsequent post-generation aggregation of per-chunk logs into per-session logs. Note that a CDN may also elect to generate aggregate per-session logs when performing HAS delivery, but this needs to be in addition to, and not instead of, the per-chunk delivery logs. We note that this may be revisited in future versions of this document.

Note that in the case of non real-time logging, the trigger of the transmission or generation of the logging file appears to be a synchronous process from a protocol standpoint. The implementation algorithm can choose to enforce a maximum size for the logging file beyond which the transmission is automatically triggered (and thus allow for an asynchronous transmission process).

2.2.2. Logging Collection

This is the process that continuously collects logs generated by the log-generating entities within a CDN.

In a CDNI environment, in addition to collecting logging information from log-generating entities within the local CDN, the Collection process also collects logging information provided by another CDN, or other CDNs, through the CDNI Logging interface. This is illustrated in Figure 2 where we see that the Collection process of the uCDN collects logging information from log-generating entities within the uCDN as well as logging information coming through CDNI Logging exchange with the dCDN through the CDNI Logging interface.

2.2.3. Logging Filtering

A CDN may require to only present different subset of the whole logging information collected to various log-consuming applications. This is achieved by the Filtering process.

In particular, the Filtering process can also filter the right subset of information that needs to be provided to a given interconnected CDN. For example, the filtering process in the dCDN can be used to ensure that only the logging information related to tasks performed on behalf of a given uCDN are made available to that uCDN (thereby filtering all the logging information related to deliveries by the dCDN of content for its own CSPs). Similarly, the Filtering process may filter or partially mask some fields, for example, to protect End Users' privacy when communicating CDNI Logging information to another CDN. Filtering of logging information prior to communication of this information to other CDNs via the CDNI Logging interface requires that the downstream CDN can recognize the set of log records that relate to each interconnected CDN.

The CDN will also filter some internal scope information such as information related to its internal alarms (security, failures, load, etc).

In some use cases described in [RFC6770], the interconnected CDNs do not want to disclose details on their internal topology. The filtering process can then also filter confidential data on the

dCDNs' topology (number of servers, location, etc.). In particular, information about the requests served by every Surrogate may be confidential. Therefore, the Logging information must be protected so that data such as Surrogates' hostnames is not disclosed to the uCDN. In the "Inter-Affiliates Interconnection" use case, this information may be disclosed to the uCDN because both the dCDN and the uCDN are operated by entities of the same group.

2.2.4. Logging Rectification and Post-Generation Aggregation

If Logging is generated periodically, it is important that the sessions that start in one Logging period and end in another are correctly reported. If they are reported in the starting period, then the Logging of this period will be available only after the end of the session, which delays the Logging generation.

A Logging rectification/update mechanism could be useful to reach a good trade-off between the Logging generation delay and the Logging accuracy. Depending on the selected Logging protocol(s), such mechanism may be invaluable for real time Logging, which must be provided rapidly and cannot wait for the end of operations in progress.

In the presence of HAS, some log-consuming applications can benefit from aggregate per-session logs. For example, for analytics, per-session logs allow display of session-related trends which are much more meaningful for some types of analysis than chunk-related trends. In the case where the log-generating entities have generated during-generation aggregate logs, those can be used by the applications. In the case where aggregate logs have not been generated, the Rectification process can be extended with a Post-Generation Aggregation process that generates per-session logs from the per-chunk logs, possibly leveraging the information included in the per-chunk logs for that purpose (Content Collection IDentifier and a Session IDentifier). However, in accordance with [RFC6983], this document does not define exchange of such aggregate logs on the CDNI Logging interface. We note that this may be revisited in future versions of this document.

2.2.5. Log-Consuming Applications

2.2.5.1. Maintenance/Debugging

Logging is useful to permit the detection (and limit the risk) of content delivery failures. In particular, Logging facilitates the detection of configuration issues.

To detect faults, Logging must enable the reporting of any CDN delivery operation success and failure. The uCDN can summarize such information into KPIs. For instance, Logging needs to allow the computation of the number of times, during a given time period, that content delivery related to a specific service succeeds/fails.

Logging enables the CDN providers to identify and troubleshoot performance degradations. In particular, Logging enables the communication of traffic data (e.g., the amount of traffic that has been forwarded by a dCDN on behalf of an uCDN over a given period of time), which is particularly useful for CDN and network planning operations.

2.2.5.2. Accounting

Logging is essential for accounting, to permit inter-CDN billing and CSP billing by uCDNs. For instance, Logging information provided by dCDNs enables the uCDN to compute the total amount of traffic delivered by every dCDN for a particular Content Provider, as well as, the associated bandwidth usage (e.g., peak, 95th percentile), and the maximum number of simultaneous sessions over a given period of time.

2.2.5.3. Analytics and Reporting

The goal of analytics is to gather any relevant information to track audience, analyze user behavior, and monitor the performance and quality of content delivery. For instance, Logging enables the CDN providers to report on content consumption (e.g., delivered sessions per content) in a specific geographic area.

The goal of reporting is to gather any relevant information to monitor the performance and quality of content delivery and allow detection of delivery issues. For instance, reporting could track the average delivery throughput experienced by End-Users in a given region for a specific CSP or content set over a period of time.

2.2.5.4. Security

The goal of security is to prevent and monitor unauthorized access, misuse, modification, and denial of access of a service. A set of information is logged for security purposes. In particular, a record of access to content is usually collected to permit the CSP to detect infringements of content delivery policies and other abnormal End User behaviors.

2.2.5.5. Legal Logging Duties

Depending on the country considered, the CDNs may have to retain specific Logging information during a legal retention period, to comply with judicial requisitions.

2.2.5.6. Notions common to multiple Log Consuming Applications

2.2.5.6.1. Logging Information Views

Within a given log-consuming application, different views may be provided to different users depending on privacy, business, and scalability constraints.

For example, an analytics tool run by the uCDN can provide one view to an uCDN operator that exploits all the logging information available to the uCDN, while the tool may provide a different view to each CSP exploiting only the logging information related to the content of the given CSP.

As another example, maintenance and debugging tools may provide different views to different CDN operators, based on their operational role.

2.2.5.6.2. Key Performance Indicators (KPIs)

This section presents, for explanatory purposes, a non-exhaustive list of Key Performance Indicators (KPIs) that can be extracted/produced from logs.

Multiple log-consuming applications, such as analytics, monitoring, and maintenance applications, often compute and track such KPIs.

In a CDNI environment, depending on the situation, these KPIs may be computed by the uCDN or by the dCDN. But it is usually the uCDN that computes KPIs, because uCDN and dCDN may have different definitions of the KPIs and the computation of some KPIs requires a vision of all the deliveries performed by the uCDN and all its dCDNs.

Here is a list of important examples of KPIs:

- o Number of delivery requests received from End-Users in a given region for each piece of content, during a given period of time (e.g., hour/day/week/month)
- o Percentage of delivery successes/failures among the aforementioned requests
- o Number of failures listed by failure type (e.g., HTTP error code) for requests received from End Users in a given region and for

each piece of content, during a given period of time (e.g., hour/day/week/month)

- o Number and cause of premature delivery termination for End Users in a given region and for each piece of content, during a given period of time (e.g., hour/day/week/month)
- o Maximum and mean number of simultaneous sessions established by End Users in a given region, for a given Content Provider, and during a given period of time (e.g., hour/day/week/month)
- o Volume of traffic delivered for sessions established by End Users in a given region, for a given Content Provider, and during a given period of time (e.g., hour/day/week/month)
- o Maximum, mean, and minimum delivery throughput for sessions established by End Users in a given region, for a given Content Provider, and during a given period of time (e.g., hour/day/week/month)
- o Cache-hit and byte-hit ratios for requests received from End Users in a given region for each piece of content, during a given period of time (e.g., hour/day/week/month)
- o Top 10 of the most popularly requested content (during a given day/week/month),
- o Terminal type (mobile, PC, STB, if this information can be acquired from the browser type header, for example).

Additional KPIs can be computed from other sources of information than the Logging, for instance, data collected by a content portal or by specific client-side application programming interfaces. Such KPIs are out of scope for the present memo.

The KPIs used depend strongly on the considered log-consuming application -- the CDN operator may be interested in different metrics than the CSP is. In particular, CDN operators are often interested in delivery and acquisition performance KPIs, information related to Surrogates' performance, caching information to evaluate the cache-hit ratio, information about the delivered file size to compute the volume of content delivered during peak hour, etc.

Some of the KPIs, for instance those providing an instantaneous vision of the active sessions for a given CSP's content, are useful essentially if they are provided in real-time. By contrast, some other KPIs, such as the one averaged on a long period of time, can be provided in non-real time.

3. CDNI Logging File

3.1. Rules

This specification uses the Augmented Backus-Naur Form (ABNF) notation and core rules of [RFC5234]. In particular, the present document uses the following rules from [RFC5234]:

CR = %x0D ; carriage return

DIGIT = %x30-39 ; 0-9

DQUOTE = %x22 ; " (Double Quote)

CRLF = CR LF ; Internet standard newline

HEXDIG = DIGIT / "A" / "B" / "C" / "D" / "E" / "F"

HTAB = %x09 ; horizontal tab

LF = %x0A ; linefeed

OCTET = %x00-FF ; 8 bits of data

The present document also uses the following rules from [RFC3986]:

host = as specified in section 3.2.2 of [RFC3986].

IPv4address = as specified in section 3.2.2 of [RFC3986].

IPv6address = as specified in section 3.2.2 of [RFC3986].

The present document also defines the following additional rules:

ADDRESS = IPv4address / IPv6address

DATE = 4DIGIT "-" 2DIGIT "-" 2DIGIT

Dates are recorded in the format YYYY-MM-DD where YYYY, MM and DD stand for the numeric year, month and day respectively. All dates are specified in Universal Time Coordinated (UTC).

DEC = 1*DIGIT ["." *DIGIT]

QSTRING = DQUOTE *NDQUOTE DQUOTE ; where

NDQUOTE = <any OCTET excluding DQUOTE> / 2DQUOTE ; whereby a DQUOTE is conveyed inside a QSTRING unambiguously by repeating it.

NHTABSTRING = *NHTAB ; where


NHTAB = <any OCTET excluding HTAB>

TIME = 2DIGIT ":" 2DIGIT ":" 2DIGIT ["." *DIGIT]

Times are recorded in the form HH:MM:SS or HH:MM:SS.S where HH is the hour in 24 hour format, MM is minutes and SS is seconds. All times are specified in Universal Time Coordinated (UTC).

3.2. CDNI Logging File Structure

As defined in Section 1.1 a CDNI logging field is as an atomic logging information element and a CDNI Logging Record is a collection of CDNI Logging Fields containing all logging information corresponding to a single logging event. This document defines a third level of structure, the CDNI Logging File, that is a collection of CDNI Logging Records. This structure is illustrated in Figure 3. The use of a file structure for transfer of CDNI Logging information is selected since this is the most common practise today for exchange of logging information within and across CDNs.



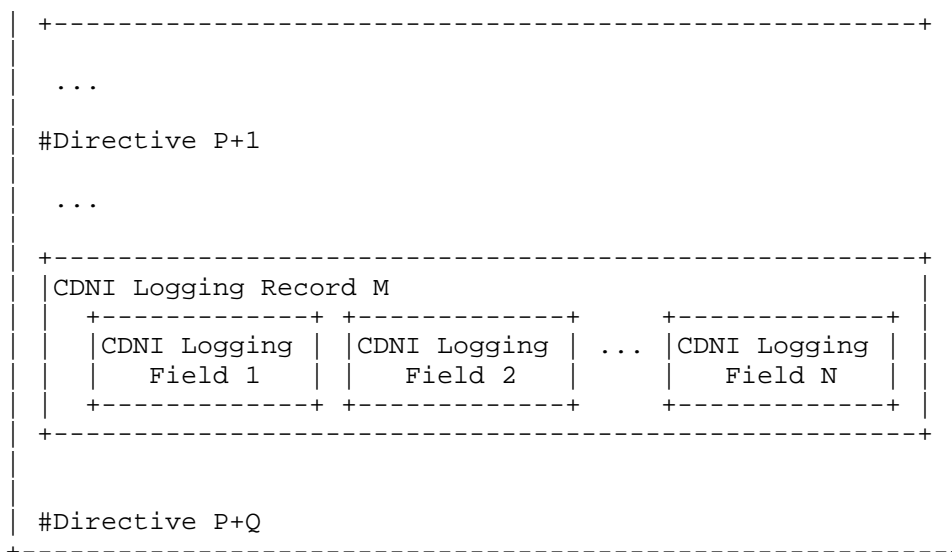


Figure 3: Structure of Logging Files

The CDNI Logging File format is inspired from the W3C Extended Log File Format [ELF]. However, it is fully specified by the present document. Where the present document differs from the W3C Extended Log File Format, an implementation of CDNI Logging MUST comply with the present document.

Using a format that resembles the W3C Extended Log File Format is intended to keep CDNI logging format close to intra-CDN logging format commonly used in CDNs today, thereby minimizing systematic translation at CDN/CDNI boundary.

A CDNI Logging File MUST contain a sequence of lines containing US-ASCII characters [CHAR_SET] terminated by CRLF.

Each line of a CDNI Logging File MUST contain either a directive or a CDNI Logging Record.

Directives record information about the CDNI Logging process itself. Lines containing directives MUST begin with the "#" character. Directives are specified in Section 3.3.

Logging Records provide actual details of the logged event. Logging Records are specified in Section 3.4.

The CDNI File structure is defined by the following rules:

```
DIRLINE = "#" directive CRLF
DIRGROUP = 1*DIRLINE
RECLINE = <CDNI Logging Record> CRLF
RECGROUP = *RECLINE
<CDNI Logging File> = 1*<DIRGROUP RECGROUP>
```

3.3. CDNI Logging File Directives

The CDNI Logging File directives are defined by the following rules:

```
directive = DIRNAME ":" HTAB DIRVAL

DIRNAME = any CDNI Logging Directive name registered in the CDNI
Logging Directive Names registry (Section 5.1).

DIRVAL = <directive value as specified below for each directive
name>
```

An implementation of the CDNI Logging interface **MUST** support all of the following directives, listed below by their directive name:

o Version:

- * format: "CDNI" "/" 1*DIGIT "." 1*DIGIT
- * directive value: indicates the version of the CDNI Logging File format. The value **MUST** be "CDNI/1.0" for the version specified in the present document.
- * occurrence: there **MUST** be one and only one instance of this directive per CDNI Logging File. It **MUST** be the first line of the CDNI Logging file.

o UUID:

- * format: NHTABSTRING
- * directive value: this a Universally Unique Identifier (UUID) from the UUID Uniform Resource Name (URN) namespace specified in [RFC4122] for the CDNI Logging File .
- * occurrence: there **MUST** be one and only one instance of this directive per CDNI Logging File.

- o Claimed-Origin:
 - * format: host
 - * directive value: this contains the claimed identification of the entity transmitting the CDNI Logging File (e.g. the host in a dCDN supporting the CDNI Logging interface) or the entity responsible for transmitting the CDNI Logging File (e.g. the dCDN).
 - * occurrence: there MUST be zero or one instance of this directive per CDNI Logging File. This directive MAY be included by the dCDN. It MUST NOT be included or modified by the uCDN.
- o Verified-Origin:
 - * format: host
 - * directive value: this contains the identification, as established by the entity receiving the CDNI Logging file, of the entity transmitting the CDNI Logging File (e.g. the host in a dCDN supporting the CDNI Logging interface) or the entity responsible for transmitting the CDNI Logging File (e.g. the dCDN).
 - * occurrence: there MUST be zero or one instance of this directive per CDNI Logging File. This directive MAY be added by the uCDN (e.g. before storing the CDNI Logging File). It MUST NOT be included by the dCDN. The mechanisms used by the uCDN to establish and validate the entity responsible for the CDNI Logging File is outside the scope of the present document. We observe that, in particular, this may be achieved through authentication mechanisms that are part of the CDNI Logging File pull mechanism (Section 4.2).
- o Record-Type:
 - * format: NHTABSTRING
 - * directive value: indicates the type of the CDNI Logging Records that follow this directive, until another Record-Type directive (or the end of the CDNI Logging File). This can be any CDNI Logging Record type registered in the CDNI Logging Record-types registry (Section 5.2). "cdni_http_request_v1" MUST be indicated as the Record-Type directive value for CDNI Logging records corresponding to HTTP request (e.g. a HTTP delivery request) as specified in Section 3.4.1.

- * occurrence: there MUST be at least one instance of this directive per CDNI Logging File. The first instance of this directive MUST precede a Fields directive and precede any CDNI Logging Record.
- o Fields:
 - * format: FIENAME *<HTAB FIENAME> ; where FIENAME can take any CDNI Logging field name registered in the CDNI Logging Field Names registry (Section 5.3).
 - * directive value: this lists the names of all the fields for which a value is to appear in the CDNI Logging Records that follow the instance of this directive (until another instance of this directive). The names of the fields, as well as their possible occurrences, are specified for each type of CDNI Logging Records in Section 3.4.
 - * occurrence: there MUST be at least one instance of this directive per Record-Type directive. The first instance of this directive for a given Record-Type MUST appear before any CDNI Logging Record for this Record-Type.
- o Integrity-Hash:
 - * format: 32HEXDIG
 - * directive value: This directive permits the detection of a corrupted CDNI Logging File. This can be useful, for instance, if a problem occurs on the filesystem of the dCDN Logging system and leads to a truncation of a logging file. The valid Integrity-Hash value is included in this directive by the entity that transmits the CDNI Logging File. It is computed by applying the MD5 ([RFC1321]) cryptographic hash function on the CDNI Logging File, including all the directives and logging records, up to the Integrity-Hash directive itself, excluding the Integrity-Hash directive itself. The Integrity-Hash value is represented as a US-ASCII encoded hexadecimal number, 32 digits long (representing a 128 bit hash value). The entity receiving the CDNI Logging File also computes in a similar way the MD5 hash on the received CDNI Logging File and compares this hash to the value of the Integrity-Hash directive. If the two values are equal, then the received CDNI Logging File MUST be considered non-corrupted. If the two values are different, the received CDNI Logging File MUST be considered corrupted. The behavior of the entity that received a corrupted CDNI Logging File is outside the scope of this specification; we note that the entity MAY attempt to pull again the same CDNI

Logging file from the transmitting entity. If the entity receiving the CDNI Logging File adds a Verified-Origin directive, it MUST recompute and update the Integrity-Hash directive so it also protects the added Verified-Origin directive.

- * occurrence: there MUST be zero or one instance of this directive. There SHOULD be one instance of this directive. One situation where that directive could be omitted is where integrity protection is already provided via another mechanism (for example if an integrity hash is associated to the CDNI Logging file out of band through the CDNI Logging Logging Feed Section 4.1 leveraging ATOM extensions such as those proposed in [I-D.snell-atompub-link-extensions]. When present, this field MUST be the last line of the CDNI Logging File.

3.4. CDNI Logging Records

A CDNI Logging Record consists of a sequence of CDNI Logging Fields relating to that single CDNI Logging Record.

CDNI Logging Fields MUST be separated by the "horizontal tabulation (HTAB)" character.

To facilitate readability, a prefix scheme is used for CDNI Logging field names in a similar way to the one used in W3C Extended Log File Format [ELF] . The semantics of the prefix in the present document is:

- o c: refers to the User Agent that issues the request (corresponds to the "client" of W3C Extended Log Format)
- o d: refers to the dCDN (relative to a given CDN acting as a uCDN)
- o s: refers to the dCDN Surrogate that serves the request (corresponds to the "server" of W3C Extended Log Format)
- o u: refers to the uCDN (relative to a given CDN acting as a dCDN)
- o cs: refers to communication from the User-Agent towards the dCDN Surrogate
- o sc: refers to communication from the dCDN Surrogate towards the User-Agent

An implementation of the CDNI Logging interface as per the present specification MUST support the CDNI HTTP Delivery Records as specified in Section 3.4.1.

A CDNI Logging Record is defined by the following rules:

FIEVAL = <CDNI Logging Field value>

<CDNI Logging Record> = FIEVAL *<HTAB FIEVAL> ; where FIEVAL contains the CDNI Logging field values corresponding to the CDNI Logging field names (FIENAME) listed in the last Fields directive preceding the present CDNI Logging Record.

3.4.1. HTTP Request Logging Record

The HTTP Request Logging Record is a CDNI Logging Record of Record-Type "cdni_http_request_v1". It contains the following CDNI Logging Fields, listed by their field name:

- o date:
 - * format: DATE
 - * field value: the date at which the processing of request completed on the Surrogate.
 - * occurrence: there MUST be one and only one instance of this field.
- o time:
 - * format: TIME
 - * field value: the time at which the processing of request completed on the Surrogate.
 - * occurrence: there MUST be one and only one instance of this field.
- o time-taken:
 - * format: DEC
 - * field value: decimal value of the duration, in seconds, between the start of the processing of the request and the completion of the request processing (e.g. completion of delivery) by the Surrogate.
 - * occurrence: there MUST be one and only one instance of this field.
- o c-ip:

- * format: ADDRESS
- * field value: the source IPv4 or IPv6 address (i.e. the "client" address) in the request received by the Surrogate.
- * occurrence: there MUST be one and only one instance of this field.
- o c-ip-anonimizing:
 - * format: 1*DIGIT
 - * field value: the number of rightmost bits of the address in the c-ip field that are zeroed-out in order to anonymize the logging record. The mechanism by which the two ends of the CDNI Logging interface agree on whether anonymization is to be supported and the number of bits that need to be zeroed-out for this purpose are outside the scope of the present document.
 - * occurrence: there MUST be zero or one instance of this field.
- o c-port:
 - * format: 1*DIGIT
 - * field value: the source TCP port (i.e. the "client" port) in the request received by the Surrogate.
 - * occurrence: there MUST be zero or exactly one instance of this field.
- o s-ip:
 - * format: ADDRESS
 - * field value: the IPv4 or IPv6 address of the Surrogate that served the request (i.e. the "server" address).
 - * occurrence: there MUST be zero or exactly one instance of this field.
- o s-hostname:
 - * format: host
 - * field value: the hostname of the Surrogate that served the request (i.e. the "server" hostname).

- * occurrence: there MUST be zero or exactly one instance of this field.
- o s-port:
 - * format: 1*DIGIT
 - * field value: the destination TCP port (i.e. the "server" port) in the request received by the Surrogate.
 - * occurrence: there MUST be zero or exactly one instance of this field.
- o cs-method:
 - * format: NHTABSTRING
 - * field value: this is the HTTP method of the HTTP request received by the Surrogate.
 - * occurrence: There MUST be one and only one instance of this field.
- o cs-uri:
 - * format: NHTABSTRING
 - * field value: this is the complete URL of the request received by the Surrogate. It is exactly in the format of a http_URL specified in [RFC2616]) or, when the request was a HTTPS request ([RFC2818]), it is in the format of a http_URL but with the scheme part set to "https" instead of "http".
 - * occurrence: there MUST be zero or exactly one instance of this field.
- o u-uri:
 - * format: NHTABSTRING
 - * field value: this is a complete URL, derived from the complete URI of the request received by the Surrogate (i.e. the cs-uri) but transformed by the entity generating or transmitting the CDNI Logging Record, in a way that is agreed upon between the two ends of the CDNI Logging interface, so the transformed URI is meaningful to the uCDN. For example, the two ends of the CDNI Logging interface could agree that the u-uri is constructed from the cs-uri by removing the part of the

hostname that exposes which individual Surrogate actually performed the delivery. The details of modification performed to generate the u-uri, as well as the mechanism to agree on these modifications between the two sides of the CDNI Logging interface are outside the scope of the present document.

- * occurrence: there MUST be one and only one instance of this field.
- o protocol:
 - * format: NHTABSTRING
 - * field value: this is value of the HTTP-Version field as specified in [RFC2616] of the Request-Line of the request received by the Surrogate (e.g. "HTTP/1.1").
 - * occurrence: there MUST be one and only one instance of this field.
- o sc-status:
 - * format: 3DIGIT
 - * field value: this is the HTTP Status-Code in the HTTP response from the Surrogate.
 - * occurrence: There MUST be one and only one instance of this field.
- o sc-total-bytes:
 - * format: 1*DIGIT
 - * field value: this is the total number of bytes of the HTTP response sent by the Surrogate in response to the request. This includes the bytes of the Status-Line (including HTTP headers) and of the message-body.
 - * occurrence: There MUST be one and only one instance of this field.
- o sc-entity-bytes:
 - * format: 1*DIGIT
 - * field value: this is the number of bytes of the message-body in the HTTP response sent by the Surrogate in response to the

request. This does not include the bytes of the Status-Line (and therefore does not include the bytes of the HTTP headers).

- * occurrence: there MUST be zero or exactly one instance of this field.
- o cs(<HTTP-header-name>):
 - * format: QSTRING
 - * field value: the value of the HTTP header (identified by the <HTTP-header-name> in the CDNI Logging field name) as it appears in the request processed by the Surrogate. For example, when the CDNI Logging field name (FIENAME) listed in the prededing Fields directive is "cs(User-Agent)", this CDNI Logging field value contains the value of the User-Agent HTTP header as received by the Surrogate in the request it processed.
 - * occurrence: there MUST be zero, one or any number of instance of this field.
- o sc(<HTTP-header-name>):
 - * format: QSTRING
 - * field value: the value of the HTTP header (identified by the <HTTP-header-name> in the CDNI Logging field name) as it appears in the response issued by the Surrogate to serve the request.
 - * occurrence: there MUST be zero, one or any number of instance of this field.
- o s-ccid:
 - * format: QSTRING
 - * field value: this contains the value of the Content Collection Identifier associated by the uCDN to the content served by the Surrogate via the CDNI Metadata interface ([I-D.ietf-cdni-metadata]).
 - * occurrence: there MUST be zero or exactly one instance of this field.
- o s-sid:

- * format: QSTRING
 - * field value: this contains the value of a Session Identifier generated by the dCDN for a specific HTTP Adaptive Streaming (HAS) session and whose value is included in the Logging record for every content chunk delivery of that session in view of facilitating the later correlation of all the per content chunk log records of a given HAS session. See section 3.4.2.2. of [RFC6983] for more discussion on the concept of Session Identifier.
 - * occurrence: there MUST be zero or exactly one instance of this field.
- o s-cached:
- * format: 1DIGIT
 - * field value: this characterises whether the Surrogate served the request using content already stored on its local cache or not. The allowed values are "0" (for miss) and "1" (for hit). "1" MUST be used when the Surrogate did serve the request using exclusively content already stored on its local cache. "0" MUST be used otherwise (including cases where the Surrogate served the request using some, but not all, content already stored on its local cache). Note that a "0" only means a cache miss in the Surrogate and does not provide any information on whether the content was already stored, or not, in another device of the dCDN i.e. whether this was a "dCDN hit" or "dCDN miss".
 - * occurrence: there MUST be zero or exactly one instance of this field.

The "Fields" directive corresponding to a HTTP Request Logging Record MUST list all the fields name whose occurrence is specified above as "There MUST be one and only one instance of this field". The corresponding fields value MUST be present in every HTTP Request Logging Record.

The "Fields" directive corresponding to a HTTP Request Logging Record MAY list all the fields value whose occurrence is specified above as "there MUST be zero or exactly one instance of this field" or "there MUST be zero, one or any number of instance of this field". The set of such fields name actually listed in the "Fields" directive is selected by the implementation generating the CDNI Logging File based on agreements between the interconnected CDNs established through mechanisms outside the scope of this specification (e.g. contractual agreements). When such a field name is not listed in the "Fields"

directive, the corresponding field value MUST NOT be included in the Logging Record. When such a field name is listed in the "Fields" directive, the corresponding field value MUST be included in the Logging Record; in that case, if the value for the field is not available, this MUST be conveyed via a dash character ("-").

The fields name listed in the "Fields" directive MAY be listed in the order in which they are listed in Section 3.4.1 or MAY be listed in any other order.

A dCDN-side implementation of the CDNI Logging interface MUST support the ability to include valid values for the following Logging Fields in a CDNI Logging Record of Record-Type "cdni_http_request_v1":

- o date
- o time
- o time-taken
- o c-ip
- o c-port
- o s-ip
- o s-hostname
- o s-port
- o cs- method
- o cs-uri
- o u-uri
- o protocol
- o sc-status
- o sc- total-bytes
- o sc-entity-bytes
- o cs(<HTTP-header>)
- o sc(<HTTP-header>)

- o s-cached

A dCDN-side implementation of the CDNI Logging interface MAY support the ability to include valid values for the following Logging Fields in a CDNI Logging Record of Record-Type "cdni_http_request_v1":

- o c-ip-anonimizing

- o s-ccid

- o s-sid

An uCDN-side implementation of the CDNI Logging interface MUST be able to accept CDNI Logging Files with CDNI Logging Records of Record-Type "cdni_http_request_v1" containing any CDNI Logging Field defined in Section 3.4.1 as long as the CDNI Logging Record and the CDNI Logging File are compliant with the present document.

3.5. CDNI Logging File Example

```
#Version:<HTAB>CDNI/1.0<CRLF>
```

```
#UUID:<HTAB>"urn:uuid:f81d4fae-7dec-11d0-a765-00a0c91e6bf6"<CRLF>
```

```
#Claimed-Origin:<HTAB>cdni-logging-entity.dcdn.example.com<CRLF>
```

```
#Record-Type:<HTAB>cdni_http_request_v1<CRLF>
```

```
#Fields:<HTAB>date<HTAB>time<HTAB>time-taken<HTAB>c-ip<HTAB>cs-  
method<HTAB>u-uri<HTAB>protocol<HTAB>sc-status<HTAB>sc-total-  
bytes<HTAB>cs(User-Agent)<HTAB>cs(Referer)<HTAB>s-cached<CRLF>
```

```
2013-05-17<HTAB>00:38:06.825<HTAB>9.058<HTAB>10.5.7.1<HTAB>GET<HTAB>h  
ttp://cdni-ucdn.dcdn.example.com/video/movie100.mp4<HTAB>HTTP/  
1.1<HTAB>200<HTAB>6729891<HTAB>"Mozilla/5.0 (Windows; U; Windows NT  
6.0; en-US) AppleWebKit/533.4 (KHTML, like Gecko) Chrome/5.0.375.127  
Safari /533.4"<HTAB>"host1.example.com"<HTAB>1<CRLF>
```

```
2013-05-17<HTAB>00:39:09.145<HTAB>15.32<HTAB>10.5.10.5<HTAB>GET<HTAB>  
http://cdni-ucdn.dcdn.example.com/video/movie118.mp4<HTAB>HTTP/  
1.1<HTAB>200<HTAB>15799210<HTAB>"Mozilla/5.0 (Windows; U; Windows NT  
6.0; en-US) AppleWebKit/533.4 (KHTML, like Gecko) Chrome/5.0.375.127  
Safari /533.4"<HTAB>"host1.example.com"<HTAB>1<CRLF>
```

```
2013-05-17<HTAB>00:42:53.437<HTAB>52.879<HTAB>10.5.10.5<HTAB>GET<HTAB>  
>http://cdni-ucdn.dcdn.example.com/video/picture11.mp4<HTAB>HTTP/  
1.0<HTAB>200<HTAB>97234724<HTAB>"Mozilla/5.0 (Windows; U; Windows NT  
6.0; en-US) AppleWebKit/533.4 (KHTML, like Gecko) Chrome/5.0.375.127  
Safari /533.4"<HTAB>"host5.example.com"<HTAB>0<CRLF>
```

```
#Integrity-Hash:<HTAB>fe113dfce8fec91323a4fc02261af26e<CRLF>
```

4. CDNI Logging File Exchange Protocol

This document specifies a protocol for the exchange of CDNI Logging Files as specified in Section 3.

This protocol comprises:

- o a CDNI Logging feed, allowing the dCDN to notify the uCDN about the CDNI Logging files that can be retrieved by that uCDN from the dCDN, as well as all the information necessary for retrieving each of these CDNI Logging File. The CDNI Logging feed is specified in Section 4.1.
- o a CDNI Logging File pull mechanism, allowing the uCDN to obtain from the dCDN a given CDNI Logging File at the uCDN convenience. The CDNI Logging File pull mechanisms is specified in Section 4.2.

An implementation of the CDNI Logging interface as per the present document generating CDNI Logging file (i.e. on the dCDN side) MUST support the server side of the CDNI Logging feed and the server side of the CDNI Logging pull mechanism.

An implementation of the CDNI Logging interface as per the present document consuming CDNI Logging file (i.e. on the uCDN side) MUST support the client side of the CDNI Logging feed and the client side of the CDNI Logging pull mechanism.

We note that implementations of the CDNI Logging interface MAY also support other mechanisms to exchange CDNI Logging Files, for example in view of exchanging logging information with minimum time-lag (e.g. sub-minute or sub-second) between when the event occurred in the dCDN and when the corresponding Logging Record is made available to the uCDN (e.g. for log-consuming applications requiring extremely fresh logging information such as near-real-time content delivery monitoring). Such mechanisms are outside the scope of the present document but might be defined in future version of this document .

4.1. CDNI Logging Feed

The server-side implementation of the CDNI Logging feed MUST produce an Atom feed [RFC4287]. This feed is used to advertise log files that are available for the client-side to retrieve using the CDNI Logging pull mechanism.

4.1.1. Atom Formatting

A CDNI Logging feed MUST be structured as an Archived feed, as defined in [RFC5005], and MUST be formatted in Atom [RFC4287]. This means it consists of a subscription document that is regularly updated as new CDNI logging files become available, and information about older CDNI Logging files is moved into archive documents. Once created, archive documents are never modified.

Each CDNI Logging file listed in an Atom feed MUST be described in an atom:entry container element.

The atom:entry MUST contain an atom:content element whose "src" attribute is a link to the CDNI Logging file and whose "type" attribute is the MIME Media Type indicating that the entry is a CDNI Logging File. We define this MIME Media Type as "application/cdni.LoggingFile" (See Section 5.4).

For compatibility with some Atom feed readers the atom:entry MAY also contain an atom:link entry whose "href" attribute is a link to the CDNI Logging file and whose "type" attribute is the MIME Media Type indicating that the entry is a CDNI Logging File using the "application/cdni.LoggingFile" MIME Media Type (See Section 5.4).

The IRI used in the atom:id of the atom:entry MUST contain the UUID of the CDNI Logging file.

The atom:updated in the atom:entry MUST indicate the time at which the CDNI Logging file was last updated.

4.1.2. Updates to Log Files and the Feed

CDNI Logging files MUST NOT be modified by the dCDN once published in the CDNI Logging feed.

The frequency with which the subscription feed is updated, the period of time covered by each CDNI Logging file or each archive document, and timeliness of publishing of CDNI Logging files are outside the scope of the present document and are expected to be agreed upon by uCDN and dCDN via other means (e.g. human agreement).

The server-side implementation SHOULD use HTTP cache control headers on the subscription feed to indicate the frequency at which the client-side is to poll for updates.

The potential retention limits (e.g. sliding time window) within which the dCDN is to retain and be ready to serve an archive document is outside the scope of the present document and is expected to be agreed upon by uCDN and dCDN via other means (e.g. human agreement). The server-side implementation MUST retain, and be ready to serve, any archive document within the agreed retention limits. Outside these agreed limits, the server-side implementation MAY be unable to serve (e.g., with HTTP status code 404) an archive document or MAY refuse to serve it (e.g., with HTTP status code 403 or 410).

4.1.3. Redundant Feeds

The server-side implementation MAY present more than one CDNI Logging feed and for redundancy, CDNI Logging files MAY be published in more than one feed.

A client-side implementation MAY support such redundant CDNI Logging feeds. If it supports redundant CDNI Logging feed, the client-side SHOULD use the UUID of the CDNI Logging file, presented in the atom:id element of the Atom feed, to avoid unnecessarily pulling and storing each CDNI Logging file more than once.

4.1.4. Example CDNI Logging Feed

Figure 4 illustrates an example of the subscription document of a CDNI Logging feed.

```
<?xml version="1.0" encoding="utf-8"?>
<feed xmlns="http://www.w3.org/2005/Atom"
<http://www.w3.org/2005/Atom%22>>
  <title type="text">CDNI Logging Feed</title>
  <updated>2013-03-23T14:46:11Z</updated>
  <id>urn:uuid:663ae677-40fb-e99a-049d-c5642916b8ce</id>
  <link href="https://dcdn.example/logfeeds/ucdn1"
    rel="self" type="application/atom+xml" />
  <link href="https://dcdn.example/logfeeds/ucdn1"
    rel="current" type="application/atom+xml" />
  <link href="https://dcdn.example/logfeeds/ucdn1/201303231400"
    rel="prev-archive" type="application/atom+xml" />
  <generator version="example version 1">CDNI Log Feed
    Generator</generator>
  <author><name>dcdn.example</name></author>
  <entry>
    <title type="text">CDNI Logging File for uCDN at
```

```

2013-03-23 14:15:00</title>
<id>urn:uuid:12345678-1234-abcd-00aa-01234567abcd</id>
<updated>2013-03-23T14:15:00Z</updated>
<content src="https://dcdn.example/logs/ucdn/
  http-requests-201303231415000000000"
  type="application/cdni.LoggingFile" />
<summary>CDNI Logging File for uCDN at
2013-03-23 14:15:00</summary>
</entry>
<entry>
  <title type="text">CDNI Logging File for uCDN at
    2013-03-23 14:30:00</title>
  <id>urn:uuid:87654321-4321-dcba-aa00-dcba7654321</id>
  <updated>2013-03-23T14:30:00Z</updated>
  <content src="https://dcdn.example/logs/ucdn/
    http-requests-201303231430000000000"
    type="application/cdni.LoggingFile" />
  <summary>CDNI Logging File for uCDN at
    2013-03-23 15:30:00</summary>
</entry>
...
<entry>
  ...
</entry>
</feed>

```

Figure 4: Example subscription document of a CDNI Logging Feed

4.2. CDNI Logging File Pull

A client-side implementation of the CDNI Logging interface MAY pull, at its convenience, a CDNI Logging File that is published by the server-side in the CDNI Logging Feed (in the subscription document or an archive document). To do so, the client-side:

- o MUST use HTTP v1.1 ([RFC2616]);
- o SHOULD use TLS (i.e. use what is loosely referred to as "HTTPS") as per [RFC2818] whenever protection of the CDNI Logging information is required (see Section 6.1);
- o MUST use the URI that was associated to the CDNI Logging File (within the "src" attribute of the corresponding atom:content element) in the CDNI Logging Feed
- o MUST support exchange of CDNI Logging Files with no content encoding applied to the representation;

- o SHOULD support exchange of CDNI Logging Files with "gzip" content encoding (as defined in [RFC2616]) applied to the representation.

Note that a client-side implementation of the CDNI Logging interface MAY pull a CDNI Logging File that it has already pulled.

The server-side implementation MUST respond to valid pull request by a client-side implementation for a CDNI Logging File published by the server-side in the CDNI Logging Feed (in the subscription document or an archive document). The server-side implementation:

- o MUST handle the client-side request as per HTTP v1.1;
- o MUST include the CDNI Logging File identified by the request URI inside the body of the HTTP response;
- o MUST support exchange of CDNI Logging Files with no content encoding applied to the representation;
- o SHOULD support exchange of CDNI Logging Files with "gzip" content encoding (as defined in [RFC2616]) applied to the representation.

Content negotiation approaches defined in [RFC2616] (e.g. using Accept-Encoding request-header field or Content-Encoding entity-header field) MAY be used by the client-side and server-side implementations to establish the content-coding to be used for a particular exchange of a CDNI Logging File.

Applying compression content encoding (such as "gzip") is expected to mitigate the impact of exchanging the large volumes of logging information expected across CDNs. This is expected to be particularly useful in the presence of HTTP Adaptive Streaming (HAS) which, as per the present version of the document, will result in a separate CDNI Log Record for each HAS segment delivery in the CDNI Logging File.

The potential retention limits (e.g. sliding time window, maximum aggregate file storage quotas) within which the dCDN is to retain and be ready to serve a CDNI Logging File previously advertised in the CDNI Logging Feed is outside the scope of the present document and is expected to be agreed upon by uCDN and dCDN via other means (e.g. human agreement). The server-side implementation MUST retain, and be ready to serve, any CDNI Logging File within the agreed retention limits. Outside these agreed limits, the server-side implementation MAY be unable to serve (e.g., with HTTP status code 404) a CDNI Logging File or MAY refuse to serve it (e.g., with HTTP status code 403 or 410).

5. IANA Considerations

5.1. CDNI Logging Directive Names Registry

The IANA is requested to create a new registry, CDNI Logging Directive Names.

The initial contents of the CDNI Logging File Directives registry comprise the names of the directives specified in Section 3.3 of the present document, and are as follows:

Directive Name	Reference
Version	+ RFC xxxx
UUID	+ RFC xxxx
Claimed-Origin	+ RFC xxxx
Verified-Origin	+ RFC xxxx
Record-Type	+ RFC xxxx
Fields	+ RFC xxxx
Integrity-Hash	+ RFC xxxx

Figure 5

[Instructions to IANA: Replace "RFC xxxx" above by the RFC number of the present document]

Within the registry, names are to be allocated by IANA according to the "Specification Required" policy specified in [RFC5226].

5.2. CDNI Logging Record-Types Registry

The IANA is requested to create a new registry, CDNI Logging Record-Types.

The initial contents of the CDNI Logging Record-Types registry comprise the names of the CDNI Logging Record types specified in Section 3.4 of the present document, and are as follows:

Record-Types	Reference
cdni_http_request_v1	+ RFC xxxx

Figure 6

[Instructions to IANA: Replace "RFC xxxx" above by the RFC number of the present document]

Within the registry, Record-Types are to be allocated by IANA according to the "Specification Required" policy specified in [RFC5226].

5.3. CDNI Logging Field Names Registry

The IANA is requested to create a new registry, CDNI Logging Field Names.

The initial contents of the CDNI Logging Fields Names registry comprise the names of the CDNI Logging fields specified in Section 3.4 of the present document, and are as follows:

Field Name	Reference
date	+ RFC xxxx
time	+ RFC xxxx
time-taken	+ RFC xxxx
c-ip	+ RFC xxxx
c-ip-anonimizing	+ RFC xxxx
c-port	+ RFC xxxx
s-ip	+ RFC xxxx
s-hostname	+ RFC xxxx
s-port	+ RFC xxxx
cs- method	+ RFC xxxx
cs-uri	+ RFC xxxx
u-uri	+ RFC xxxx
protocol	+ RFC xxxx
sc-status	+ RFC xxxx
sc- total-bytes	+ RFC xxxx
sc-entity-bytes	+ RFC xxxx
cs(<HTTP-header>)	+ RFC xxxx
sc(<HTTP-header>)	+ RFC xxxx
s-ccid	+ RFC xxxx
s-sid	+ RFC xxxx
s-cached	+ RFC xxxx

Figure 7

[Instructions to IANA: Replace "RFC xxxx" above by the RFC number of the present document]

Within the registry, names are to be allocated by IANA according to the "Specification Required" policy specified in [RFC5226].

5.4. CDNI Logging MIME Media Type

The IANA is requested to allocate the "application/cdni.LoggingFile" MIME Media Type (whose use is specified in Section 4.1.1 of the present document) in the MIME Media Types registry.

6. Security Considerations

6.1. Authentication, Confidentiality, Integrity Protection

The use of TLS as per [RFC2818] for transport of the CDNI Logging feed mechanism (Section 4.1) and CDNI Logging File pull mechanism (Section 4.2) allows:

- o the dCDN and uCDN to authenticate each other (to ensure they are transmitting/receiving CDNI Logging File from an authenticated CDN)
- o the CDNI Logging information to be transmitted with confidentiality
- o the integrity of the CDNI Logging information to be protected during the exchange

In an environment where any such protection is required, TLS SHOULD be used for transport of the CDNI Logging feed and the CDNI Logging File pull.

A CDNI Logging implementation MUST support TLS transport of the CDNI Logging feed and the CDNI Logging File pull.

The Integrity-Hash directive inside the CDNI Logging File provides additional integrity protection, this time targeting potential corruption of the CDNI logging information during the CDNI Logging File generation. This mechanism does not allow restoration of the corrupted CDNI Logging information, but it allows detection of such corruption and therefore triggering of appropriate correcting actions (e.g. discard of corrupted information, attempt to re-obtain the CDNI Logging information).

6.2. Denial of Service

This document does not define specific mechanism to protect against Denial of Service (DoS) attacks on the Logging Interface. However, the CDNI Logging feed and CDNI Logging pull endpoints can be protected against DoS attacks through the use of TLS transport and/or via mechanisms outside the scope of the CDNI Logging interface such as firewalling or use of Virtual Private Networks (VPNs).

Protection of dCDN Surrogates against spoofed delivery requests is outside the scope of the CDNI Logging interface.

6.3. Privacy

CDNs have the opportunity to collect detailed information about the downloads performed by End-Users. The provision of this information to another CDN introduces potential End-Users privacy protection concerns. We observe that when CDNI interconnection is realised as per [I-D.ietf-cdni-framework], the uCDN handles the initial End-User requests (before it is redirected to the dCDN) so, regardless of which information is, or is not, communicated to the uCDN through the CDNI Logging interface, the uCDN has visibility on significant information such as the IP address of the End-User request and the URL of the request. Nonetheless, if the dCDN and uCDN agree that anonymization is required to avoid making some detailed information available to the uCDN (such as how much bytes of the content has been watched by an enduser and/or at what time) or is required to meet some legal obligations, then the uCDN and dCDN can agree to exchange anonymized End-User IP addresses in CDNI Logging files and the c-ip-anonymization field can be used to convey the number of bits that have been anonymized so that the meaningful information can still be easily extracted from the anonymized addressses (e.g. for geolocation aware analytics).

7. Acknowledgments

This document borrows from the W3C Extended Log Format [ELF].

Rob Murray significantly contributed into the text of Section 4.1 .

The authors thank Ben Niven-Jenkins, Kevin Ma, David Mandelberg and Ray van Brandenburg for their ongoing input.

Finally, we also thank Sebastien Cubaud, Pawel Grochocki, Christian Jacquenet, Yannick Le Louedec, Anne Marrec , Emile Stephan, Fabio Costa, Sara Oueslati, Yvan Massot, Renaud Edel, Joel Favier and the contributors of the EU FP7 OCEAN project for their input in the early versions of this document.

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2616, June 1999.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005.
- [RFC4122] Leach, P., Mealling, M., and R. Salz, "A Universally Unique Identifier (UUID) URN Namespace", RFC 4122, July 2005.
- [RFC4287] Nottingham, M., Ed. and R. Sayre, Ed., "The Atom Syndication Format", RFC 4287, December 2005.
- [RFC5005] Nottingham, M., "Feed Paging and Archiving", RFC 5005, September 2007.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [RFC5234] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, January 2008.

8.2. Informative References

- [CHAR_SET] , "IANA Character Sets registry", , <<http://www.iana.org/assignments/character-sets/character-sets.xml>>.
- [ELF] Phillip M. Hallam-Baker, . and . Brian Behlendorf, "Extended Log File Format, W3C (work in progress), WD-logfile-960323", , <<http://www.w3.org/TR/WD-logfile.html>>.
- [I-D.ietf-cdni-framework] Peterson, L. and B. Davie, "Framework for CDN Interconnection", draft-ietf-cdni-framework-05 (work in progress), September 2013.
- [I-D.ietf-cdni-metadata]

Niven-Jenkins, B., Murray, R., Watson, G., Caulfield, M., Leung, K., and K. Ma, "CDN Interconnect Metadata", draft-ietf-cdni-metadata-02 (work in progress), July 2013.

[I-D.ietf-cdni-requirements]

Leung, K. and Y. Lee, "Content Distribution Network Interconnection (CDNI) Requirements", draft-ietf-cdni-requirements-10 (work in progress), September 2013.

[I-D.snell-atompub-link-extensions]

Snell, J., "Atom Link Extensions", draft-snell-atompub-link-extensions-09 (work in progress), June 2012.

[RFC1321] Rivest, R., "The MD5 Message-Digest Algorithm", RFC 1321, April 1992.

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

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

[RFC6770] Bertrand, G., Stephan, E., Burbridge, T., Eardley, P., Ma, K., and G. Watson, "Use Cases for Content Delivery Network Interconnection", RFC 6770, November 2012.

[RFC6983] van Brandenburg, R., van Deventer, O., Le Faucheur, F., and K. Leung, "Models for HTTP-Adaptive-Streaming-Aware Content Distribution Network Interconnection (CDNI)", RFC 6983, July 2013.

Appendix A. Compliance with CDNI Requirements

[Editor's Note: This appendix is intended to help the WG understand compliance of the CDNI Logging interface against the requirements defined in the CDNI requirements document, in order to establish readiness for the document publication. This appendix is expected to be removed for publication].

[Editor's Note: this appendix may need a small update if ietf-cdni-requirements introduces an additional requirement for Privacy/Anonimization as recently discussed on the list, and if LI14 & LI-15 are modified]

The three tables below review compliance against, respectively, the Generic CDNI requirements, the CDNI Logging interface requirements and the CDNI security requirements of [I-D.ietf-cdni-requirements]. The first two columns of the tables indicate the requirement number,

and the requirement priority as defined in [I-D.ietf-cdni-requirements]. The third column of the table indicates the level of compliance of the CDNI Logging interface specified in the present document against the given requirement, and the fourth column provides additional comment and explanation on how or why the compliance is achieved or not achieved.

Re- quire- ment	Prior- ity	Compli- ance	Comment
GEN-1	MED	Full	Leverages existing protocols incl including HTTP, TLS and ATOM
GEN-2	HIGH	Full	Does not require any change or upgrade to the user agent
GEN-3	HIGH	Full	Does not require any change or upgrade to the Content Service Provider
GEN-4	HIGH	Full	Does not depend on intra-CDN info
GEN-5	HIGH	Full	Supports logging of HTTP delivery
GEN-6	HIGH	N/A	
GEN-7	LOW	Not Compliant	Only supports logging for HTTP delivery, but easily extensible to add support for other delivery protos
GEN-8	LOW	N/A	
GEN-9	MED	Full	Supports logging across cascaded CDNs
GEN-10	MED	Full	Supports any topology of interconnected CDNs
GEN-11	HIGH	Partial	No explicit mechanism for loop avoidance is defined; the exchange of logs is usually done in a point to point manner between two well identified entities situated in the uCDN and dCDN. Loop avoidance is expected to be handled by implementations based on inferring the CDN path from the URI structure in the HTTP redirection case and/or administrative information

			(topology restrictions in case of DNS redirection method also handled administratively)
GEN-12	HIGH	N/A	
GEN-13	HIGH	Full	Supports Logging for HTTP Adaptive Streaming (HAS) content, with one Logging Record per HAS segment. Supports a few optional logging fields specific to HAS. Does not support summarized Logging Records for HAS, but extensible to add that.

Figure 8: Compliance to Generic CDNI Requirements

Re-requirement	Priority	Compliance	Comment
LI-1	HIGH	Full	Reliable transfer is achieved by the transport protocol: the logging data is transmitted over HTTP over TCP. Also, supports optional redundancy of the Logging feed.
LI-2	HIGH	Full	Supports logs for all content deliveries both complete and incomplete performed by the dCDN on behalf of the uCDN
LI-3	MED	Full	The CDNI Logging Interface does not impose any restrictions related to the transmission of logs generated by intermediary CDNs; the dCDN formats internally all the final logging files including those received from intermediary CDNs and the locally generated
LI-4	HIGH	Full	The ATOM feed allows the uCDN to trigger the download of logging files whenever needed
LI-5	MED	Partial	The uCDN can pull logging files from the dCDN whenever a new file is available. The timing constraints for

			the generation of the logging files are to be defined offline, and can be defined to an arbitrary period. This is expected to be compatible with applications that have low timing constraints (e.g. 24 hours) such as billing. This is expected to be compatible with applications that have high timing constraints (e.g. 5 minutes) such as monitoring or analytics. This is not expected to be compatible with applications that have very high timing constraints (e.g. a few seconds or below)
LI-6	HIGH	Full	Section 3.4 describes the CDNI Logging Records and the possible fields that can be included in a record. Supports a single type of CDNI event i.e. HTTP delivery
LI-7	HIGH	Full	Defines an ATOM based feed and HTTP or HTTPS transport
LI-8	MED	Partial	Allows as uCDN to pull current CDNI Logging files to access current Logging records. Does not allow uCDN to request Log Records before next Logging file is made available.
LI-9	LOW	Not Compliant	The current version of the document does not specify any mechanisms for producing aggregate / summarized logs, but exchanged logging files provide all the information that is necessary to the uCDN in order to produce aggregated logs. Extensible to add such mechanisms in the future
LI-10	LOW	Not compliant	Future versions might define such a mechanism for logging performance data. Allows uCDN to derive some perf indicators from delivery Records
LI-11	MED	Not compliant	Future versions might define such a mechanism for logging data about resources consumed by the dCDN

LI-12	MED	Not compliant	Future versions might define such a mechanism for logging data about resources consumed by cascaded CDNs
LI-13	HIGH	Not compliant	Not supported by CDNI Logging interface. However, it is expected that the CDNI Control interface will allow tracing of delete request results (e.g. success, failure).
LI-14	HIGH	Full	Details about extensibility mechanisms in Section 6.
LI-15	HIGH	Full	Details about proprietary fields in Section 6.
LI-16	HIGH	Full	The CDNI Logging feed indicates which Logging file is (or was) available
LI-17	MED	Full	Content Collection ID and Session ID are supported for logging records related to HTTP Adaptive Streaming

Figure 9: Compliance to CDNI Logging interface Requirements

Requirement	Priority	Compliance	Comment
SEC-1	HIGH	Full	TLS can be used for transport of any CDNI logging related information which provides authentication, confidentiality, integrity protection as well as protection against spoofing and replay
SEC-2	HIGH	Partial	No specific mechanism against Denial of Service attacks is defined on the Logging Interface. Spoofed requests can be avoided by using TLS. Protection against spoofed delivery requests are outside the scope of CDNI Logging.
SEC-3	MED	N/A	Establishing CDN path with non-repudiation is outside the scope of CDNI Logging. Does not prevent use of

			such mechanism (e.g. including info in content URI).
SEC-4	MED	Not compliant	A non-repudiation mechanism for CDNI logging might be defined in a separate document
SEC-5	LOW	N/A	

Figure 10: Compliance to CDNI Security Requirements

Authors' Addresses

Francois Le Faucheur (editor)
Cisco Systems
E.Space Park - Batiment D
6254 Allee des Ormes - BP 1200
Mougins cedex 06254
FR

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

Gilles Bertrand (editor)
Orange
38-40 rue du General Leclerc
Issy les Moulineaux 92130
FR

Phone: +33 1 45 29 89 46
Email: gilles.bertrand@orange.com

Iuniana Oprescu (editor)
Orange
38-40 rue du General Leclerc
Issy les Moulineaux 92130
FR

Phone: +33 6 89 06 92 72
Email: iuniana.oprescu@orange.com

Roy Peterkofsky
Skytide, Inc.
One Kaiser Plaza, Suite 785
Oakland CA 94612
USA

Phone: +01 510 250 4284
Email: roy@skytide.com

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: December 10, 2016

F. Le Faucheur, Ed.

G. Bertrand, Ed.

I. Oprescu, Ed.

R. Peterkofsky
Google Inc.
June 8, 2016

CDNI Logging Interface
draft-ietf-cdni-logging-27

Abstract

This memo specifies the Logging interface between a downstream CDN (dCDN) and an upstream CDN (uCDN) that are interconnected as per the CDN Interconnection (CDNI) framework. First, it describes a reference model for CDNI logging. Then, it specifies the CDNI Logging File format and the actual protocol for exchange of CDNI Logging Files.

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 December 10, 2016.

Copyright Notice

Copyright (c) 2016 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
1.2. Requirements Language	5
2. CDNI Logging Reference Model	5
2.1. CDNI Logging interactions	5
2.2. Overall Logging Chain	8
2.2.1. Logging Generation and During-Generation Aggregation	9
2.2.2. Logging Collection	10
2.2.3. Logging Filtering	10
2.2.4. Logging Rectification and Post-Generation Aggregation	11
2.2.5. Log-Consuming Applications	12
2.2.5.1. Maintenance/Debugging	12
2.2.5.2. Accounting	13
2.2.5.3. Analytics and Reporting	13
2.2.5.4. Content Protection	13
2.2.5.5. Notions common to multiple Log Consuming Applications	14
3. CDNI Logging File	16
3.1. Rules	16
3.2. CDNI Logging File Structure	17
3.3. CDNI Logging Directives	20
3.4. CDNI Logging Records	24
3.4.1. HTTP Request Logging Record	25
3.5. CDNI Logging File Extension	36
3.6. CDNI Logging File Examples	36
3.7. Cascaded CDNI Logging Files Example	39
4. Protocol for Exchange of CDNI Logging File After Full Collection	42
4.1. CDNI Logging Feed	43
4.1.1. Atom Formatting	43
4.1.2. Updates to Log Files and the Feed	43
4.1.3. Redundant Feeds	44
4.1.4. Example CDNI Logging Feed	44
4.2. CDNI Logging File Pull	46
5. Protocol for Exchange of CDNI Logging File During Collection	47
6. IANA Considerations	48
6.1. CDNI Logging Directive Names Registry	48
6.2. CDNI Logging File version Registry	48
6.3. CDNI Logging record-types Registry	49

6.4. CDNI Logging Field Names Registry	50
6.5. CDNI Logging MIME Media Type	51
7. Security Considerations	52
7.1. Authentication, Authorization, Confidentiality, Integrity Protection	52
7.2. Denial of Service	53
7.3. Privacy	53
8. Acknowledgments	55
9. References	55
9.1. Normative References	55
9.2. Informative References	57
Authors' Addresses	59

1. Introduction

This memo specifies the CDNI Logging interface between a downstream CDN (dCDN) and an upstream CDN (uCDN). First, it describes a reference model for CDNI logging. Then, it specifies the CDNI Logging File format and the actual protocol for exchange of CDNI Logging Files.

The reader should be familiar with the following documents:

- o CDNI problem statement [RFC6707] and framework [RFC7336] identify a Logging interface,
- o Section 8 of [RFC7337] specifies a set of requirements for Logging,
- o [RFC6770] outlines real world use-cases for interconnecting CDNs. These use cases require the exchange of Logging information between the dCDN and the uCDN.

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

The present document describes:

- o The CDNI Logging reference model (Section 2),
- o The CDNI Logging File format (Section 3),
- o The CDNI Logging File Exchange protocol (Section 4).

1.1. Terminology

In this document, the first letter of each CDNI-specific term is capitalized. We adopt the terminology described in [RFC6707] and [RFC7336], and extend it with the additional terms defined below.

Intra-CDN Logging information: logging information generated and collected within a CDN. The format of the Intra-CDN Logging information may be different to the format of the CDNI Logging information.

CDNI Logging information: logging information exchanged across CDNs using the CDNI Logging Interface.

Logging information: logging information generated and collected within a CDN or obtained from another CDN using the CDNI Logging Interface.

CDNI Logging Field: an atomic element of information that can be included in a CDNI Logging Record. The time an event/task started, the IP address of an End User to whom content was delivered, and the Uniform Resource Identifier (URI) of the content delivered, are examples of CDNI Logging fields.

CDNI Logging Record: an information record providing information about a specific event. This comprises a collection of CDNI Logging fields.

CDNI Logging File: a file containing CDNI Logging Records, as well as additional information facilitating the processing of the CDNI Logging Records.

CDN Reporting: the process of providing the relevant information that will be used to create a formatted content delivery report provided to the CSP in deferred time. Such information typically includes aggregated data that can cover a large period of time (e.g., from hours to several months). Uses of Reporting include the collection of charging data related to CDN services and the computation of Key Performance Indicators (KPIs).

CDN Monitoring: the process of providing or displaying content delivery information in a timely fashion with respect to the corresponding deliveries. Monitoring typically includes visibility of the deliveries in progress for service operation purposes. It presents a view of the global health of the services as well as information on usage and performance, for network services supervision and operation management. In particular, monitoring data can be used to generate alarms.

1.2. Requirements Language

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

2. CDNI Logging Reference Model

2.1. CDNI Logging interactions

The CDNI logging reference model between a given uCDN and a given dCDN involves the following interactions:

- o customization by the uCDN of the CDNI Logging information to be provided by the dCDN to the uCDN (e.g., control of which CDNI Logging fields are to be communicated to the uCDN for a given task performed by the dCDN or control of which types of events are to be logged). The dCDN takes into account this CDNI Logging customization information to determine what Logging information to provide to the uCDN, but it may, or may not, take into account this CDNI Logging customization information to influence what CDN logging information is to be generated and collected within the dCDN (e.g., even if the uCDN requests a restricted subset of the logging information, the dCDN may elect to generate a broader set of logging information). The mechanism to support the customization by the uCDN of CDNI Logging information is outside the scope of this document and left for further study. Until such a mechanism is available, the uCDN and dCDN are expected to agree off-line on what exact set of CDNI Logging information is to be provided by the dCDN to the uCDN, and to rely on management plane actions to configure the CDNI Logging functions in the dCDN to generate this information set and in the uCDN to expect this information set.
- o generation and collection by the dCDN of the intra-CDN Logging information related to the completion of any task performed by the dCDN on behalf of the uCDN (e.g., delivery of the content to an End User) or related to events happening in the dCDN that are relevant to the uCDN (e.g., failures or unavailability in dCDN). This takes place within the dCDN and does not directly involve CDNI interfaces.
- o communication by the dCDN to the uCDN of the Logging information collected by the dCDN relevant to the uCDN. This is supported by the CDNI Logging interface and in the scope of the present document. For example, the uCDN may use this Logging information to charge the CSP, to perform analytics and monitoring for

operational reasons, to provide analytics and monitoring views on its content delivery to the CSP or to perform trouble-shooting. This document exclusively specifies non-real-time exchange of Logging information. Closer to real-time exchange of Logging information (say sub-minute or sub-second) is outside the scope of the present document and left for further study. This document exclusively specifies exchange of Logging information related to content delivery. Exchange of Logging information related to operational events (e.g., dCDN request routing function unavailable, content acquisition failure by dCDN) for audit or operational reactive adjustments by uCDN is outside the scope of the present document and left for further study.


- o customization by the dCDN of the CDNI Logging information to be provided by the uCDN on behalf of the dCDN. The mechanism to support the customization by the dCDN of CDNI Logging information is outside the scope of this document and left for further study.
- o generation and collection by the uCDN of Intra-CDN Logging information related to the completion of any task performed by the uCDN on behalf of the dCDN (e.g., serving of content by uCDN to dCDN for acquisition purposes by dCDN) or related to events happening in the uCDN that are relevant to the dCDN. This takes place within the uCDN and does not directly involve CDNI interfaces.
- o communication by the uCDN to the dCDN of the Logging information collected by the uCDN relevant to the dCDN. For example, the dCDN might potentially benefit from this information for security auditing or content acquisition troubleshooting. This is outside the scope of this document and left for further study.

Figure 1 provides an example of CDNI Logging interactions (focusing only on the interactions that are in the scope of this document) in a particular scenario where four CDNs are involved in the delivery of content from a given CSP: the uCDN has a CDNI interconnection with dCDN-1 and dCDN-2. In turn, dCDN-2 has a CDNI interconnection with dCDN-3, where dCDN-2 is acting as an upstream CDN relative to dCDN-3. In this example, uCDN, dCDN-1, dCDN-2 and dCDN-3 all participate in the delivery of content for the CSP. In this example, the CDNI Logging interface enables the uCDN to obtain Logging information from all the dCDNs involved in the delivery. In the example, the uCDN uses the Logging information:

- o to analyze the performance of the delivery performed by the dCDNs and to adjust its operations after the fact (e.g., request routing) as appropriate,

- o to provide (non-real-time) reporting and monitoring information to the CSP.

For instance, the uCDN merges Logging information, extracts relevant KPIs, and presents a formatted report to the CSP, in addition to a bill for the content delivered by uCDN itself or by its dCDNs on the CSP's behalf. The uCDN may also provide Logging information as raw log files to the CSP, so that the CSP can use its own logging analysis tools.



==> CDNI Logging Interface
 ***> outside the scope of CDNI

Figure 1: Interactions in CDNI Logging Reference Model

A downstream CDN relative to uCDN (e.g., dCDN-2) integrates the relevant Logging information obtained from its own downstream CDNs (i.e., dCDN-3) in the Logging information that it provides to the uCDN, so that the uCDN ultimately obtains all Logging information relevant to a CSP for which it acts as the authoritative CDN. Such aggregation is further discussed in Section 3.7.

Note that the format of Logging information that a CDN provides over the CDNI interface might be different from the one that the CDN uses internally. In this case, the CDN needs to reformat the Logging information before it provides this information to the other CDN over

the CDNI Logging interface. Similarly, a CDN might reformat the Logging information that it receives over the CDNI Logging interface before injecting it into its log-consuming applications or before providing some of this Logging information to the CSP. Such reformatting operations introduce latency in the logging distribution chain and introduce a processing burden. Therefore, there are benefits in specifying CDNI Logging formats that are suitable for use inside CDNs and also are close to the intra-CDN Logging formats commonly used in CDNs today.

2.2. Overall Logging Chain

This section discusses the overall logging chain within and across CDNs to clarify how CDN Logging information is expected to fit in this overall chain. Figure 2 illustrates the overall logging chain within the dCDN, across CDNs using the CDNI Logging interface and within the uCDN. Note that the logging chain illustrated in the Figure is obviously only an example and varies depending on the specific environments. For example, there may be more or fewer instantiations of each entity (e.g., there may be 4 Log consuming applications in a given CDN). As another example, there may be one instance of Rectification process per Log Consuming Application instead of a shared one.

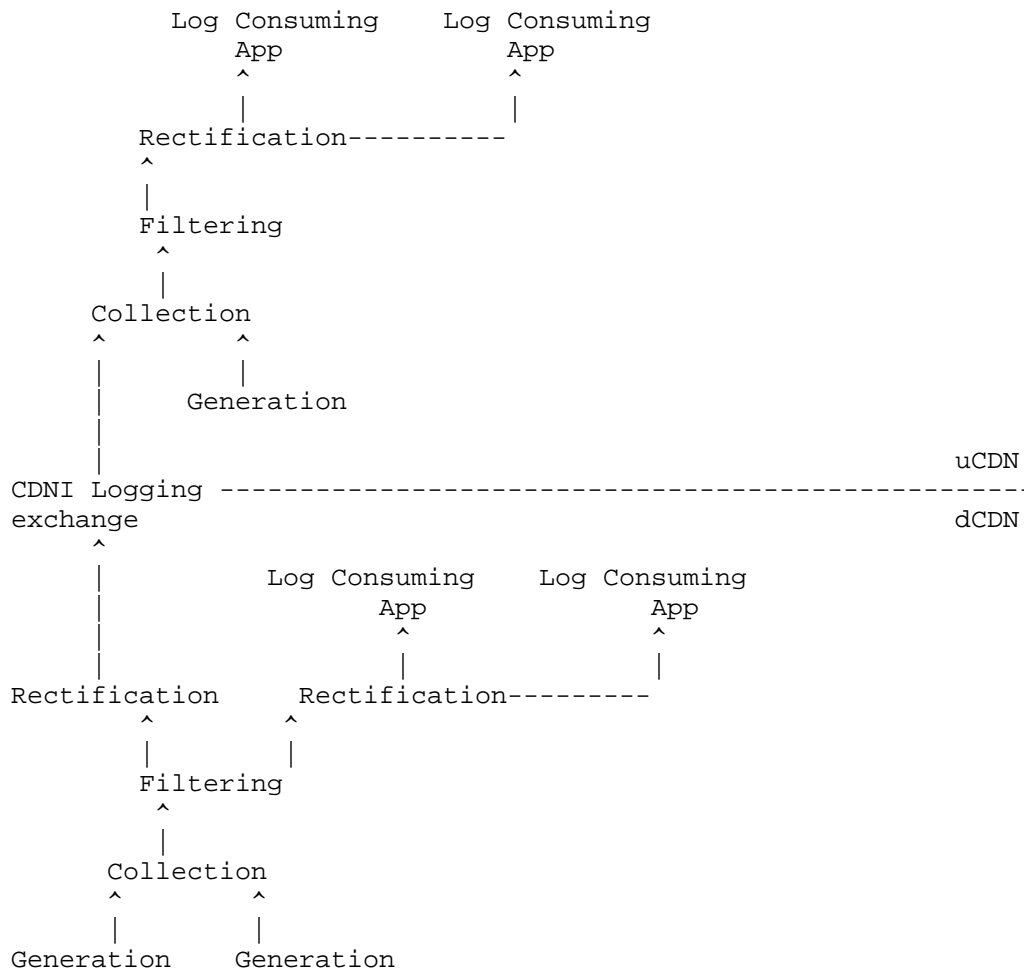


Figure 2: CDNI Logging in the overall Logging Chain

The following subsections describe each of the processes potentially involved in the logging chain of Figure 2.

2.2.1. Logging Generation and During-Generation Aggregation

CDNs typically generate Logging information for all significant task completions, events, and failures. Logging information is typically generated by many devices in the CDN including the surrogates, the request routing system, and the control system.

The amount of Logging information generated can be huge. Therefore, during contract negotiations, interconnected CDNs often agree on a retention duration for Logging information, and/or potentially on a maximum volume of Logging information that the dCDN ought to keep. If this volume is exceeded, the dCDN is expected to alert the uCDN but may not keep more Logging information for the considered time period. In addition, CDNs may aggregate Logging information and transmit only summaries for some categories of operations instead of the full Logging information. Note that such aggregation leads to an information loss, which may be problematic for some usages of the Logging information (e.g., debugging).

[RFC6983] discusses logging for HTTP Adaptive Streaming (HAS). In accordance with the recommendations articulated there, it is expected that a surrogate will generate separate Logging information for delivery of each chunk of HAS content. This ensures that separate Logging information can then be provided to interconnected CDNs over the CDNI Logging interface. Still in line with the recommendations of [RFC6983], the Logging information for per-chunk delivery may include some information (a Content Collection Identifier and a Session Identifier) intended to facilitate subsequent post-generation aggregation of per-chunk logs into per-session logs. Note that a CDN may also elect to generate aggregate per-session logs when performing HAS delivery, but this needs to be in addition to, and not instead of, the per-chunk delivery logs. We note that aggregate per-session logs for HAS delivery are for further study and outside the scope of this document.

2.2.2. Logging Collection

This is the process that continuously collects Logging information generated by the log-generating entities within a CDN.

In a CDNI environment, in addition to collecting Logging information from log-generating entities within the local CDN, the Collection process also collects Logging information provided by another CDN, or other CDNs, through the CDNI Logging interface. This is illustrated in Figure 2 where we see that the Collection process of the uCDN collects Logging information from log-generating entities within the uCDN as well as Logging information coming from the dCDNs through the CDNI Logging interface.

2.2.3. Logging Filtering

A CDN may be required to only present different subsets of the whole Logging information collected to various log-consuming applications. This is achieved by the Filtering process.

In particular, the Filtering process can also filter the right subset of Logging information that needs to be provided to a given interconnected CDN. For example, the filtering process in the dCDN can be used to ensure that only the Logging information related to tasks performed on behalf of a given uCDN are made available to that uCDN (thereby filtering out all the Logging information related to deliveries by the dCDN of content for its own CSPs). Similarly, the Filtering process may filter or partially mask some fields, for example, to protect End Users' privacy when communicating CDNI Logging information to another CDN. Filtering of Logging information prior to communication of this information to other CDNs via the CDNI Logging interface requires that the downstream CDN can recognize the subset of Logging information that relate to each interconnected CDN.

The CDN will also filter some internal scope information such as information related to its internal alarms (security, failures, load, etc).

In some use cases described in [RFC6770], the interconnected CDNs do not want to disclose details on their internal topology. The filtering process can then also filter confidential data on the dCDNs' topology (number of servers, location, etc.). In particular, information about the requests served by each Surrogate may be confidential. Therefore, the Logging information needs to be protected so that data such as Surrogates' hostnames are not disclosed to the uCDN. In the "Inter-Affiliates Interconnection" use case, this information may be disclosed to the uCDN because both the dCDN and the uCDN are operated by entities of the same group.

2.2.4. Logging Rectification and Post-Generation Aggregation

If Logging information is generated periodically, it is important that the sessions that start in one Logging period and end in another are correctly reported. If they are reported in the starting period, then the Logging information of this period will be available only after the end of the session, which delays the Logging information generation. A simple approach is to provide the complete Logging Record for a session in the Logging Period of the session end.

A Logging rectification/update mechanism could be useful to reach a good trade-off between the Logging information generation delay and the Logging information accuracy.

In the presence of HAS, some log-consuming applications can benefit from aggregate per-session logs. For example, for analytics, per-session logs allow display of session-related trends which are much more meaningful for some types of analysis than chunk-related trends. In the case where aggregate logs have been generated directly by the

log-generating entities, those can be used by the applications. In the case where aggregate logs have not been generated, the Rectification process can be extended with a Post-Generation Aggregation process that generates per-session logs from the per-chunk logs, possibly leveraging the information included in the per-chunk logs for that purpose (Content Collection IDentifier and a Session IDentifier). However, in accordance with [RFC6983], this document does not define exchange of such aggregate logs on the CDNI Logging interface. We note that this is for further study and outside the scope of this document.

2.2.5. Log-Consuming Applications

2.2.5.1. Maintenance/Debugging

Logging information is useful to permit the detection (and limit the risk) of content delivery failures. In particular, Logging information facilitates the detection of configuration issues.

To detect faults, Logging information needs to report success and failure of CDN delivery operations. The uCDN can summarize such information into KPIs. For instance, Logging information needs to allow the computation of the number of times, during a given time period, that content delivery related to a specific service succeeds/fails.

Logging information enables the CDN providers to identify and troubleshoot performance degradations. In particular, Logging information enables tracking of traffic data (e.g., the amount of traffic that has been forwarded by a dCDN on behalf of an uCDN over a given period of time), which is particularly useful for CDN and network planning operations.

Some of these maintenance and debugging applications only require aggregate logging information highly compatible with use of anonymization of IP addresses (as supported by the present document and specified in the definition of the c-groupid field under Section 3.4.1). However, in some situations, it may be useful, where compatible with privacy protection, to access some CDNI Logging Records containing full non-anonymized IP addresses. This is allowed in the definition of the c-groupid (under Section 3.4.1), with very significant privacy protection limitations that are discussed in the definition of the c-groupid field. For example, this may be useful for detailed fault tracking of a particular end user content delivery issue. Where there is a hard requirement by uCDN or CSP to associate a given enduser to individual CDNI Logging Records (e.g., to allow a-posteriori analysis of individual delivery for example in situations of performance-based penalties), instead of using

aggregates containing a single client as discussed in the c-groupid field definition, an alternate approach is to ensure that a client identifier is embedded in the request fields that can be logged in a CDNI Logging Record (for example by including the client identifier in the URI query string or in a HTTP Header). That latter approach offers two strong benefits: first, the aggregate inside the c-groupid can contain more than one client, thereby ensuring stronger privacy protection; second, it allows a reliable identification of the client while IP address does not in many situations (e.g., behind NAT, where dynamic IP addresses are used and reused,...). However, care SHOULD be taken that the client identifiers exposed in other fields of the CDNI Records cannot themselves be linked back to actual users.

2.2.5.2. Accounting

Logging information is essential for accounting, to permit inter-CDN billing and CSP billing by uCDNs. For instance, Logging information provided by dCDNs enables the uCDN to compute the total amount of traffic delivered by every dCDN for a particular Content Provider, as well as, the associated bandwidth usage (e.g., peak, 95th percentile), and the maximum number of simultaneous sessions over a given period of time.

2.2.5.3. Analytics and Reporting

The goals of analytics include gathering any relevant information in order to be able to develop statistics on content download, analyze user behavior, and monitor the performance and quality of content delivery. For instance, Logging information enables the CDN providers to report on content consumption (e.g., delivered sessions per content) in a specific geographic area.

The goal of reporting is to gather any relevant information to monitor the performance and quality of content delivery and allow detection of delivery issues. For instance, reporting could track the average delivery throughput experienced by End Users in a given region for a specific CSP or content set over a period of time.

2.2.5.4. Content Protection

The goal of content protection is to prevent and monitor unauthorized access, misuse, modification, and denial of access to a content. A set of information is logged in a CDN for security purposes. In particular, a record of access to content is usually collected to permit the CSP to detect infringements of content delivery policies and other abnormal End User behaviors.

2.2.5.5. Notions common to multiple Log Consuming Applications

2.2.5.5.1. Logging Information Views

Within a given log-consuming application, different views may be provided to different users depending on privacy, business, and scalability constraints.

For example, an analytics tool run by the uCDN can provide one view to an uCDN operator that exploits all the Logging information available to the uCDN, while the tool may provide a different view to each CSP exploiting only the Logging information related to the content of the given CSP.

As another example, maintenance and debugging tools may provide different views to different CDN operators, based on their operational role.

2.2.5.5.2. Key Performance Indicators (KPIs)

This section presents, for explanatory purposes, a non-exhaustive list of Key Performance Indicators (KPIs) that can be extracted/produced from logs.

Multiple log-consuming applications, such as analytics, monitoring, and maintenance applications, often compute and track such KPIs.

In a CDNI environment, depending on the situation, these KPIs may be computed by the uCDN or by the dCDN. But it is usually the uCDN that computes KPIs, because the uCDN and dCDN may have different definitions of the KPIs and the computation of some KPIs requires a vision of all the deliveries performed by the uCDN and all its dCDNs.

Here is a list of important examples of KPIs:

- o Number of delivery requests received from End Users in a given region for each piece of content, during a given period of time (e.g., hour/day/week/month)
- o Percentage of delivery successes/failures among the aforementioned requests
- o Number of failures listed by failure type (e.g., HTTP error code) for requests received from End Users in a given region and for each piece of content, during a given period of time (e.g., hour/day/week/month)

- o Number and cause of premature delivery termination for End Users in a given region and for each piece of content, during a given period of time (e.g., hour/day/week/month)
- o Maximum and mean number of simultaneous sessions established by End Users in a given region, for a given Content Provider, and during a given period of time (e.g., hour/day/week/month)
- o Volume of traffic delivered for sessions established by End Users in a given region, for a given Content Provider, and during a given period of time (e.g., hour/day/week/month)
- o Maximum, mean, and minimum delivery throughput for sessions established by End Users in a given region, for a given Content Provider, and during a given period of time (e.g., hour/day/week/month)
- o Cache-hit and byte-hit ratios for requests received from End Users in a given region for each piece of content, during a given period of time (e.g., hour/day/week/month)
- o Top 10 most popularly requested contents (during a given day/week/month)
- o Terminal type (mobile, PC, STB, if this information can be acquired from the browser type inferred from the User Agent string, for example).

Additional KPIs can be computed from other sources of information than the Logging information, for instance, data collected by a content portal or by specific client-side application programming interfaces. Such KPIs are out of scope for the present document.

The KPIs used depend strongly on the considered log-consuming application -- the CDN operator may be interested in different metrics than the CSP is. In particular, CDN operators are often interested in delivery and acquisition performance KPIs, information related to Surrogates' performance, caching information to evaluate the cache-hit ratio, information about the delivered file size to compute the volume of content delivered during peak hour, etc.

Some of the KPIs, for instance those providing an instantaneous vision of the active sessions for a given CSP's content, are useful essentially if they are provided in a timely manner. By contrast, some other KPIs, such as those averaged on a long period of time, can be provided in non-real-time.

3. CDNI Logging File

3.1. Rules

This specification uses the Augmented Backus-Naur Form (ABNF) notation and core rules of [RFC5234]. In particular, the present document uses the following rules from [RFC5234]:

CR = %x0D ; carriage return

ALPHA = %x41-5A / %x61-7A ; A-Z / a-z

DIGIT = %x30-39 ; 0-9

DQUOTE = %x22 ; " (Double Quote)

CRLF = CR LF ; Internet standard newline

HEXDIG = DIGIT / "A" / "B" / "C" / "D" / "E" / "F"

HTAB = %x09 ; horizontal tab

LF = %x0A ; linefeed

VCHAR = %x21-7E ; visible (printing) characters

OCTET = %x00-FF ; 8 bits of data

The present document also uses the following rules from [RFC3986]:

host = as specified in section 3.2.2 of [RFC3986].

IPv4address = as specified in section 3.2.2 of [RFC3986].

IPv6address = as specified in section 3.2.2 of [RFC3986].

partial-time = as specified in [RFC3339].

The present document also defines the following additional rules:

ADDRESS = IPv4address / IPv6address

ALPHANUM = ALPHA / DIGIT

DATE = 4DIGIT "-" 2DIGIT "-" 2DIGIT

; Dates are encoded as "full-date" specified in [RFC3339].

DEC = 1*DIGIT ["." 1*DIGIT]

NAMEFORMAT = ALPHANUM *(ALPHANUM / "_" / "-")

QSTRING = DQUOTE *(NDQUOTE / PCT-ENCODED) DQUOTE

NDQUOTE = %x20-21 / %x23-24 / %x26-7E / UTF8-2 / UTF8-3 / UTF8-4

; whereby a DQUOTE is conveyed inside a QSTRING unambiguously
by escaping it with PCT-ENCODED.

PCT-ENCODED = "%" HEXDIG HEXDIG

; percent encoding is used for escaping octets that might be
possible in HTTP headers such as bare CR, bare LF, CR LF, HTAB,
SP or null. These octets are rendered with percent encoding in
ABNF as specified by [RFC3986] in order to avoid considering
them as separators for the logging records.

NHTABSTRING = 1*(SP / VCHAR)

TIME = partial-time

USER-COMMENT = * (SP / VCHAR / UTF8-2 / UTF8-3 / UTF8-4)

3.2. CDNI Logging File Structure

As defined in Section 1.1: a CDNI Logging Field is as an atomic logging information element, a CDNI Logging Record is a collection of CDNI Logging fields containing all logging information corresponding to a single logging event, and a CDNI Logging File contains a collection of CDNI Logging Records. This structure is illustrated in Figure 3. The use of a file structure for transfer of CDNI Logging information is selected since this is the most common practise today for exchange of logging information within and across CDNs.




Figure 3: Structure of Logging Files

The CDNI Logging File format is inspired from the W3C Extended Log File Format [ELF]. However, it is fully specified by the present document. Where the present document differs from the W3C Extended

Log File Format, an implementation of the CDNI Logging interface MUST comply with the present document. The W3C Extended Log File Format was used as a starting point, reused where possible and expanded when necessary.

Using a format that resembles the W3C Extended Log File Format is intended to keep CDNI logging format close to the intra-CDN Logging information format commonly used in CDNs today, thereby minimizing systematic translation at CDN/CDNI boundary.

A CDNI Logging File MUST contain a sequence of lines containing US-ASCII characters [CHAR_SET] terminated by CRLF. Each line of a CDNI Logging File MUST contain either a directive or a CDNI Logging Record.

Directives record information about the CDNI Logging process itself. Lines containing directives MUST begin with the "#" character. Directives are specified in Section 3.3.

Logging Records provide actual details of the logged event. Logging Records are specified in Section 3.4.

The CDNI Logging File has a specific structure. It always starts with a directive line and the first directive it contains MUST be the version.

The directive lines form together a group that contains at least one directive line. Each directives group is followed by a group of logging records. The records group contains zero or more actual logging record lines about the event being logged. A record line consists of the values corresponding to all or a subset of the possible Logging fields defined within the scope of the record-type directive. These values MUST appear in the order defined by the fields directive.

Note that future extensions MUST be compliant with the previous description. The following examples depict the structure of a CDNILOGFILE as defined currently by the record-type "cdni_http_request_v1."

DIRLINE = "#" directive CRLF

DIRGROUP = 1*DIRLINE

RECLINE = <any subset of record values that match what is expected according to the fields directive within the immediately preceding DIRGROUP>

```
RECGROUP = *RECLINE
```

```
CDNILOGFILE = 1*(DIRGROUP RECGROUP)
```

3.3. CDNI Logging Directives

A CDNI Logging directive line contains the directive name followed by ":" HTAB and the directive value.

Directive names MUST be of the format NAMEFORMAT. All directive names MUST be registered in the CDNI Logging Directives Names registry. Directive names are case-insensitive as per the basic ABNF([RFC5234]). Unknown directives MUST be ignored. Directive values can have various formats. All possible directive values for the record-type "cdni_http_request_v1" are further detailed in this section.

The following example shows the structure of a directive and enumerates strictly the directive values presently defined in the version "cdni/1.0" of the CDNI Logging File.

```
directive = DIRNAME ":" HTAB DIRVAL
```

```
DIRNAME = NAMEFORMAT
```

```
DIRVAL = NHTABSTRING / QSTRING / host / USER-COMMENT / FIENAME *  
(HTAB FIENAME) / 64HEXDIG
```

An implementation of the CDNI Logging interface MUST support all of the following directives, listed below by their directive name:

- o version:

- * format: NHTABSTRING

- * directive value: indicates the version of the CDNI Logging File format. The entity transmitting a CDNI Logging File as per the present document MUST set the value to "cdni/1.0". In the future, other versions of CDNI Logging File might be specified; those would use a value different to "cdni/1.0" allowing the entity receiving the CDNI Logging File to identify the corresponding version. CDNI Logging File versions are case-insensitive as per the basic ABNF([RFC5234]).

- * occurrence: there MUST be one and only one instance of this directive per CDNI Logging File. It MUST be the first line of the CDNI Logging File.

- * example: "version: HTAB cdni/1.0".
- o UUID:
 - * format: NHTABSTRING
 - * directive value: this a Uniform Resource Name (URN) from the Universally Unique Identifier (UUID) URN namespace specified in [RFC4122]). The UUID contained in the URN uniquely identifies the CDNI Logging File.
 - * occurrence: there MUST be one and only one instance of this directive per CDNI Logging File.
 - * example: "UUID: HTAB NHTABSTRING".
- o claimed-origin:
 - * format: host
 - * directive value: this contains the claimed identification of the entity transmitting the CDNI Logging File (e.g., the host in a dCDN supporting the CDNI Logging interface) or the entity responsible for transmitting the CDNI Logging File (e.g., the dCDN).
 - * occurrence: there MUST be zero or exactly one instance of this directive per CDNI Logging File. This directive MAY be included by the dCDN. It MUST NOT be included or modified by the uCDN.
 - * example: "claimed-origin: HTAB host".
- o established-origin:
 - * format: host
 - * directive value: this contains the identification, as established by the entity receiving the CDNI Logging File, of the entity transmitting the CDNI Logging File (e.g., the host in a dCDN supporting the CDNI Logging interface) or the entity responsible for transmitting the CDNI Logging File (e.g., the dCDN).
 - * occurrence: there MUST be zero or exactly one instance of this directive per CDNI Logging File. This directive MAY be added by the uCDN (e.g., before storing the CDNI Logging File). It MUST NOT be included by the dCDN. The mechanisms used by the

uCDN to establish and validate the entity responsible for the CDNI Logging File is outside the scope of the present document. We observe that, in particular, this may be achieved through authentication mechanisms that are part of the transport layer of the CDNI Logging File pull mechanism (Section 4.2).

- * ABNF example: "established-origin: HTAB host".
- o remark:
 - * format: USER-COMMENT
 - * directive value: this contains comment information. Data contained in this field is to be ignored by analysis tools.
 - * occurrence: there MAY be zero, one or any number of instance of this directive per CDNI Logging File.
 - * example: "remark: HTAB USER-COMMENT".
- o record-type:
 - * format: NAMEFORMAT
 - * directive value: indicates the type of the CDNI Logging Records that follow this directive, until another record-type directive (or the end of the CDNI Logging File). This can be any CDNI Logging Record type registered in the CDNI Logging Record-types registry (Section 6.3). For example this may be "cdni_http_request_v1" as specified in Section 3.4.1. CDNI Logging record-types are case-insensitive as per the basic ABNF([RFC5234]).
 - * occurrence: there MUST be at least one instance of this directive per CDNI Logging File. The first instance of this directive MUST precede a fields directive and MUST precede all CDNI Logging Records.
 - * example: "record-type: HTAB cdni_http_request_v1".
- o fields:
 - * format: FIENAME *(HTAB FIENAME) ; where FIENAME can take any CDNI Logging field name registered in the CDNI Logging Field Names registry (Section 6.4) that is valid for the record type specified in the record-type directive.

- * directive value: this lists the names of all the fields for which a value is to appear in the CDNI Logging Records that follow the instance of this directive (until another instance of this directive). The names of the fields, as well as their occurrences, MUST comply with the corresponding rules specified in the document referenced in the CDNI Logging Record-types registry (Section 6.3) for the corresponding CDNI Logging record-type.
 - * occurrence: there MUST be at least one instance of this directive per record-type directive. The first instance of this directive for a given record-type MUST appear before any CDNI Logging Record for this record-type. One situation where more than one instance of the fields directive can appear within a given CDNI Logging File, is when there is a change, in the middle of a fairly large logging period, in the agreement between the uCDN and the dCDN about the set of fields that are to be exchanged. The multiple occurrences allow records with the old set of fields and records with the new set of fields to be carried inside the same Logging File.
 - * example: "fields: HTAB FIENAME * (HTAB FIENAME)".
- o SHA256-hash:
- * format: 64HEXDIG
 - * directive value: This directive permits the detection of a corrupted CDNI Logging File. This can be useful, for instance, if a problem occurs on the filesystem of the dCDN Logging system and leads to a truncation of a logging file. The valid SHA256-hash value is included in this directive by the entity that transmits the CDNI Logging File. It MUST be computed by applying the SHA-256 ([RFC6234]) cryptographic hash function on the CDNI Logging File, including all the directives and logging records, up to the SHA256-hash directive itself, excluding the SHA256-hash directive itself. The SHA256-hash value MUST be represented as a US-ASCII encoded hexadecimal number, 64 digits long (representing a 256 bit hash value). The entity receiving the CDNI Logging File also computes in a similar way the SHA-256 hash on the received CDNI Logging File and compares this hash to the value of the SHA256-hash directive. If the two values are equal, then the received CDNI Logging File is to be considered non-corrupted. If the two values are different, the received CDNI Logging File is to be considered corrupted. The behavior of the entity that received a corrupted CDNI Logging File is outside the scope of this specification; we note that the entity MAY attempt to pull again the same CDNI

Logging File from the transmitting entity. If the entity receiving a non-corrupted CDNI Logging File adds an established-origin directive, it MUST then recompute and update the SHA256-hash directive so it also protects the added established-origin directive.

- * occurrence: there MUST be zero or exactly one instance of this directive. There SHOULD be exactly one instance of this directive. One situation where that directive could be omitted is where integrity protection is already provided via another mechanism (for example if an integrity hash is associated to the CDNI Logging File out-of-band through the CDNI Logging Feed (Section 4.1) leveraging ATOM extensions such as those proposed in [I-D.snell-atompub-link-extensions]. When present, the SHA256-hash field MUST be the last line of the CDNI Logging File.
- * example: "SHA256-hash: HTAB 64HEXDIG".

An uCDN-side implementation of the CDNI Logging interface MUST ignore a CDNI Logging File that does not comply with the occurrences specified above for each and every directive. For example, an uCDN-side implementation of the CDNI Logging interface receiving a CDNI Logging file with zero occurrence of the version directive, or with two occurrences of the SHA256-hash, MUST ignore this CDNI Logging File.

An entity receiving a CDNI Logging File with a value set to "cdni/1.0" MUST process the CDNI Logging File as per the present document. An entity receiving a CDNI Logging File with a value set to a different value MUST process the CDNI Logging File as per the specification referenced in the CDNI Logging File version registry (see Section 6.1) if the implementation supports this specification and MUST ignore the CDNI Logging File otherwise.

3.4. CDNI Logging Records

A CDNI Logging Record consists of a sequence of CDNI Logging fields relating to that single CDNI Logging Record.

CDNI Logging fields MUST be separated by the "horizontal tabulation (HTAB)" character.

To facilitate readability, a prefix scheme is used for CDNI Logging field names in a similar way to the one used in W3C Extended Log File Format [ELF]. The semantics of the prefix in the present document is:

- o "c-" refers to the User Agent that issues the request (corresponds to the "client" of W3C Extended Log Format)
- o "d-" refers to the dCDN (relative to a given CDN acting as an uCDN)
- o "s-" refers to the dCDN Surrogate that serves the request (corresponds to the "server" of W3C Extended Log Format)
- o "u-" refers to the uCDN (relative to a given CDN acting as a dCDN)
- o "cs-" refers to communication from the User Agent towards the dCDN Surrogate
- o "sc-" refers to communication from the dCDN Surrogate towards the User Agent

An implementation of the CDNI Logging interface as per the present specification MUST support the CDNI HTTP Request Logging Record as specified in Section 3.4.1.

A CDNI Logging Record contains the corresponding values for the fields that are enumerated in the last fields directive before the current log line. Note that the order in which the field values appear is dictated by the order of the fields names in the fields directive. There SHOULD be no dependency between the various fields values.

3.4.1. HTTP Request Logging Record

This section defines the CDNI Logging Record of record-type "cdni_http_request_v1". It is applicable to content delivery performed by the dCDN using HTTP/1.0([RFC1945]), HTTP/1.1([RFC7230],[RFC7231],[RFC7232],[RFC7233],[RFC7234],[RFC7235]) or HTTPS ([RFC2818],[RFC7230]). We observe that, in the case of HTTPS delivery, there may be value in logging additional information specific to the operation of HTTP over TLS and we note that this is outside the scope of the present document and may be addressed in a future document defining another CDNI Logging Record or another version of the HTTP Request Logging Record.

The "cdni_http_request_v1" record-type is also expected to be applicable to HTTP/2 [RFC7540] since a fundamental design tenet of HTTP/2 is to preserve the HTTP/1.1 semantics. We observe that, in the case of HTTP/2 delivery, there may be value in logging additional information specific to the additional functionality of HTTP/2 (e.g., information related to connection identification, to stream identification, to stream priority and to flow control). We note

that such additional information is outside the scope of the present document and may be addressed in a future document defining another CDNI Logging Record or another version of the HTTP Request Logging Record.

The "cdni_http_request_v1" record-type contains the following CDNI Logging fields, listed by their field name:

- o date:
 - * format: DATE
 - * field value: the date at which the processing of request completed on the Surrogate.
 - * occurrence: there MUST be one and only one instance of this field.
- o time:
 - * format: TIME
 - * field value: the time, which MUST be expressed in Coordinated Universal Time (UTC), at which the processing of request completed on the Surrogate.
 - * occurrence: there MUST be one and only one instance of this field.
- o time-taken:
 - * format: DEC
 - * field value: decimal value of the duration, in seconds, between the start of the processing of the request and the completion of the request processing (e.g., completion of delivery) by the Surrogate.
 - * occurrence: there MUST be one and only one instance of this field.
- o c-groupid:
 - * format: NHTABSTRING
 - * field value: an opaque identifier for an aggregate set of clients, derived from the client IPv4 or IPv6 address in the request received by the Surrogate and/or other network-level

identifying information. The c-groupid serves to group clients into aggregates. Example aggregates include civil geolocation information (the country, second-level administrative division, or postal code from which the client is presumed to make the request based on a geolocation database lookup) or network topological information (e.g., the BGP AS number announcing the prefix containing the address). The c-groupid MAY be structured e.g., US/TN/MEM/38138. Agreement between the dCDN and the uCDN on a mapping between IPv4 and IPv6 addresses and aggregates is presumed to occur out-of-band. The aggregation mapping SHOULD be chosen such that each aggregate contains more than one client.

- + When the aggregate is chosen so that it contains a single client (e.g., to allow more detailed analytics, or to allow a-posteriori analysis of individual delivery for example in situations of performance-based penalties) the c-groupid MAY be structured where some elements identify aggregates and one element identifies the client, e.g., US/TN/MEM/38138/43a5bdd6-95c4-4d62-be65-7410df0021e2. In the case where the aggregate is chosen so that it contains a single client:
 - the element identifying the client SHOULD be algorithmically generated (from the client IPv4 or IPv6 address in the request received by the Surrogate and/or other network-level identifying information) in a way that SHOULD NOT be linkable back to the global addressing context and that SHOULD vary over time (to offer protection against long term attacks).
 - It is RECOMMENDED that the mapping varies at least once every 24 hours.
 - The algorithmic mapping and variation over time can, in some cases, allow the uCDN (with the knowledge of the algorithm and time variation and associated attributes and keys) to reconstruct the actual client IPv4 or IPv6 address and/or other network-level identifying information when required (e.g., to allow a-posteriori analysis of individual delivery for example in situations of performance-based penalties). However, these enduser addresses SHOULD only be reconstructed on-demand and the CDNI Logging File SHOULD only be stored with the anonymised c-groupid value.
 - Allowing reconstruction of client address information carries with it grave risks to end-user privacy. Since

the c-groupid is in this case equivalent in identification power to a client IP address, its use may be restricted by regulation or law as personally identifiable information. For this reason, such use is NOT RECOMMENDED.

- One method for mapping that MAY be supported by implementations relies on a symmetric key that is known only to the uCDN and dCDN and HMAC-based Extract-and-Expand Key Derivation Function (HKDF) key derivation ([RFC5869]), as will be used in TLS 1.3 ([I-D.ietf-tls-rfc5246-bis]). When that method is used:
 - o The uCDN and dCDN need to agree on the "salt" and "input keying material", as described in Section 2.2 of [RFC5869] and the initial "info" parameter (which could be something like the business names of the two organizations in UTF-8, concatenated), as described in Section 2.3 of [RFC5869]. The hash SHOULD be either SHA-2 or SHA-3 [SHA-3] and the encryption algorithm SHOULD be 128-bit AES [AES] in Galois Counter Mode (GCM) [GCM] (AES-GCM) or better. The PRK SHOULD be chosen by both parties contributing alternate random bytes until sufficient length exists. After the initial setup, client-information can be encrypted using the key generated by the "expand" step of Section 2.3 of [RFC5869]. The encrypted value SHOULD be hex encoded or base64 encoded (as specified in section 4 of [RFC4648]). At the agreed-upon expiration time, a new key SHOULD be generated and used. New keys SHOULD be indicated by prefixing the key with a special character such as exclamation point. In this way, shorter lifetimes can be used as needed.
- * occurrence: there MUST be one and only one instance of this field.
- o s-ip:
 - * format: ADDRESS
 - * field value: the IPv4 or IPv6 address of the Surrogate that served the request (i.e., the "server" address).
 - * occurrence: there MUST be zero or exactly one instance of this field.

- o s-hostname:
 - * format: host
 - * field value: the hostname of the Surrogate that served the request (i.e., the "server" hostname).
 - * occurrence: there MUST be zero or exactly one instance of this field.
- o s-port:
 - * format: 1*DIGIT
 - * field value: the destination TCP port (i.e., the "server" port) in the request received by the Surrogate.
 - * occurrence: there MUST be zero or exactly one instance of this field.
- o cs-method:
 - * format: NHTABSTRING
 - * field value: this is the method of the request received by the Surrogate. In the case of HTTP delivery, this is the HTTP method in the request.
 - * occurrence: There MUST be one and only one instance of this field.
- o cs-uri:
 - * format: NHTABSTRING
 - * field value: this is the "effective request URI" of the request received by the Surrogate as specified in [RFC7230]. It complies with the "http" URI scheme or the "https" URI scheme as specified in [RFC7230]). Note that cs-uri can be privacy sensitive. In that case, and where appropriate, u-uri could be used instead of cs-uri.
 - * occurrence: there MUST be zero or exactly one instance of this field.
- o u-uri:
 - * format: NHTABSTRING

- * field value: this is a complete URI, derived from the "effective request URI" ([RFC7230]) of the request received by the Surrogate (i.e., the cs-uri) but transformed by the entity generating or transmitting the CDNI Logging Record, in a way that is agreed upon between the two ends of the CDNI Logging interface, so the transformed URI is meaningful to the uCDN. For example, the two ends of the CDNI Logging interface could agree that the u-uri is constructed from the cs-uri by removing the part of the hostname that exposes which individual Surrogate actually performed the delivery. The details of modification performed to generate the u-uri, as well as the mechanism to agree on these modifications between the two sides of the CDNI Logging interface are outside the scope of the present document.
- * occurrence: there MUST be one and only one instance of this field.
- o protocol:
 - * format: NHTABSTRING
 - * field value: this is value of the HTTP-Version field as specified in [RFC7230] of the Request-Line of the request received by the Surrogate (e.g., "HTTP/1.1").
 - * occurrence: there MUST be one and only one instance of this field.
- o sc-status:
 - * format: 3DIGIT
 - * field value: this is the Status-Code in the response from the Surrogate. In the case of HTTP delivery, this is the HTTP Status-Code in the HTTP response.
 - * occurrence: There MUST be one and only one instance of this field.
- o sc-total-bytes:
 - * format: 1*DIGIT
 - * field value: this is the total number of bytes of the response sent by the Surrogate in response to the request. In the case of HTTP delivery, this includes the bytes of the Status-Line,

the bytes of the HTTP headers and the bytes of the message-body.

- * occurrence: There MUST be one and only one instance of this field.
- o sc-entity-bytes:
 - * format: 1*DIGIT
 - * field value: this is the number of bytes of the message-body in the HTTP response sent by the Surrogate in response to the request. This does not include the bytes of the Status-Line or the bytes of the HTTP headers.
 - * occurrence: there MUST be zero or exactly one instance of this field.
- o cs(insert_HTTP_header_name_here):
 - * format: QSTRING
 - * field value: the value of the HTTP header (identified by the insert_HTTP_header_name_here in the CDNI Logging field name) as it appears in the request processed by the Surrogate, but prepended by a DQUOTE and appended by a DQUOTE. For example, when the CDNI Logging field name (FIENAME) listed in the preceding fields directive is cs(User-Agent), this CDNI Logging field value contains the value of the User-Agent HTTP header as received by the Surrogate in the request it processed, but prepended by a DQUOTE and appended by a DQUOTE. If the HTTP header as it appeared in the request processed by the Surrogate contains one or more DQUOTE, each DQUOTE MUST be escaped with percent encoding. For example, if the HTTP header contains My_Header"value", then the field value of the cs(insert_HTTP_header_name_here) is "My_Header%x22value%x22". The entity transmitting the CDNI Logging File MUST ensure that the respective insert_HTTP_header_name_here of the cs(insert_HTTP_header_name_here) listed in the fields directive comply with HTTP specifications. In particular, this field name does not include any HTAB, since this would prevent proper parsing of the fields directive by the entity receiving the CDNI Logging File.
 - * occurrence: there MAY be zero, one or any number of instance of this field.
- o sc(insert_HTTP_header_name_here):

- * format: QSTRING
 - * field value: the value of the HTTP header (identified by the `insert_HTTP_header_name_here` in the CDNI Logging field name) as it appears in the response issued by the Surrogate to serve the request, but prepended by a DQUOTE and appended by a DQUOTE. If the HTTP header as it appeared in the request processed by the Surrogate contains one or more DQUOTE, each DQUOTE MUST be escaped with percent encoding. For example, if the HTTP header contains `My_Header"value"`, then the field value of the `sc(insert_HTTP_header_name_here)` is `"My_Header%x22value%x22"`. The entity transmitting the CDNI Logging File MUST ensure that the respective `insert_HTTP_header_name_here` of the `cs(insert_HTTP_header_name_here)` listed in the fields directive comply with HTTP specifications. In particular, this field name does not include any HTAB, since this would prevent proper parsing of the fields directive by the entity receiving the CDNI Logging File.
 - * occurrence: there MAY be zero, one or any number of instances of this field. For a given `insert_HTTP_header_name_here`, there MUST be zero or exactly one instance of this field.
- o s-ccid:
 - * format: QSTRING
 - * field value: this contains the value of the Content Collection Identifier (CCID) associated by the uCDN to the content served by the Surrogate via the CDNI Metadata interface ([I-D.ietf-cdni-metadata]), prepended by a DQUOTE and appended by a DQUOTE. If the CCID conveyed in the CDNI Metadata interface contains one or more DQUOTE, each DQUOTE MUST be escaped with percent encoding. For example, if the CCID conveyed in the CDNI Metadata interface is `My_CCIDD"value"`, then the field value of the s-ccid is `"My_CCID%x22value%X22"`.
 - * occurrence: there MUST be zero or exactly one instance of this field. For a given `insert_HTTP_header_name_here`, there MUST be zero or exactly one instance of this field.
 - o s-sid:
 - * format: QSTRING
 - * field value: this contains the value of a Session Identifier (SID) generated by the dCDN for a specific HTTP session, prepended by a DQUOTE and appended by a DQUOTE. In particular,

for HTTP Adaptive Streaming (HAS) session, the Session Identifier value is included in the Logging record for every content chunk delivery of that session in view of facilitating the later correlation of all the per content chunk log records of a given HAS session. See section 3.4.2.2. of [RFC6983] for more discussion on the concept of Session Identifier in the context of HAS. If the SID conveyed contains one or more DQUOTE, each DQUOTE MUST be escaped with percent encoding. For example, if the SID is My_SID"value", then the field value of the s-sid is "My_SID%x22value%x22".

- * occurrence: there MUST be zero or exactly one instance of this field.
- o s-cached:
 - * format: 1DIGIT
 - * field value: this characterises whether the Surrogate served the request using content already stored on its local cache or not. The allowed values are "0" (for miss) and "1" (for hit). "1" MUST be used when the Surrogate did serve the request using exclusively content already stored on its local cache. "0" MUST be used otherwise (including cases where the Surrogate served the request using some, but not all, content already stored on its local cache). Note that a "0" only means a cache miss in the Surrogate and does not provide any information on whether the content was already stored, or not, in another device of the dCDN, i.e., whether this was a "dCDN hit" or "dCDN miss".
 - * occurrence: there MUST be zero or exactly one instance of this field.

CDNI Logging field names are case-insensitive as per the basic ABNF([RFC5234]). The "fields" directive corresponding to a HTTP Request Logging Record MUST contain all the fields names whose occurrence is specified above as "There MUST be one and only one instance of this field". The corresponding fields value MUST be present in every HTTP Request Logging Record.

The "fields" directive corresponding to a HTTP Request Logging Record MAY list all the fields value whose occurrence is specified above as "there MUST be zero or exactly one instance of this field" or "there MAY be zero, one or any number of instances of this field". The set of such field names actually listed in the "fields" directive is selected by the CDN generating the CDNI Logging File based on agreements between the interconnected CDNs established through mechanisms outside the scope of this specification (e.g., contractual

agreements). When such a field name is not listed in the "fields" directive, the corresponding field value MUST NOT be included in the Logging Record. When such a field name is listed in the "fields" directive, the corresponding field value MUST be included in the Logging Record; if the value for the field is not available, this MUST be conveyed via a dash character ("-").

The fields names listed in the "fields" directive MAY be listed in the order in which they are listed in Section 3.4.1 or MAY be listed in any other order.

Logging some specific fields from HTTP requests and responses can introduce serious security and privacy risks. For example, cookies will often contain (months) long lived token values that can be used to log into a service as the relevant user. Similar values may be included in other header fields or within URLs or elsewhere in HTTP requests and responses. Centralising such values in a CDNI Logging File can therefore represent a significant increase in risk both for the user and the web service provider, but also for the CDNs involved. Implementations ought therefore to attempt to lower the probability of such bad outcomes e.g. by only allowing a configured set of headers to be added to CDNI Logging Records, or by not supporting wildcard selection of HTTP request/response fields to add. Such mechanisms can reduce the probability that security (or privacy) sensitive values are centralised in CDNI Logging Files. Also, when agreeing on which HTTP request/response fields are to be provided in CDNI Logging Files, the uCDN and dCDN administrators ought to consider these risks. Furthermore, CDNs making use of c-groupid to identify an aggregate of clients rather than individual clients ought to realize that by logging certain header fields they may create the possibility to re-identify individual clients. In these cases heeding the above advice, or not logging header fields at all, is particularly important if the goal is to provide logs that do not identify individual clients."

A dCDN-side implementation of the CDNI Logging interface MUST implement all the following Logging fields in a CDNI Logging Record of record-type "cdni_http_request_v1", and MUST support the ability to include valid values for each of them:

- o date
- o time
- o time-taken
- o c-groupid

- o s-ip
- o s-hostname
- o s-port
- o cs-method
- o cs-uri
- o u-uri
- o protocol
- o sc-status
- o sc-total-bytes
- o sc-entity-bytes
- o cs(insert_HTTP_header_name_here)
- o sc(insert_HTTP_header_name_here)
- o s-cached

A dCDN-side implementation of the CDNI Logging interface MAY support the following Logging fields in a CDNI Logging Record of record-type "cdni_http_request_v1":

- o s-ccid
- o s-sid

If a dCDN-side implementation of the CDNI Logging interface supports these fields, it MUST support the ability to include valid values for them.

An uCDN-side implementation of the CDNI Logging interface MUST be able to accept CDNI Logging Files with CDNI Logging Records of record-type "cdni_http_request_v1" containing any CDNI Logging Field defined in Section 3.4.1 as long as the CDNI Logging Record and the CDNI Logging File are compliant with the present document.

In case an uCDN-side implementation of the CDNI Logging interface receives a CDNI Logging File with HTTP Request Logging Records that do not contain field values for exactly the set of field names actually listed in the preceding "fields" directive, the

implementation MUST ignore those HTTP Request Logging Records, and MUST accept the other HTTP Request Logging Records.

To ensure that the logging file is correct, the text MUST be sanitized before being logged. Null, bare CR, bare LF and HTAB have to be removed by escaping them through percent encoding to avoid confusion with the logging record separators.

3.5. CDNI Logging File Extension

The CDNI Logging File contains blocks of directives and blocks of corresponding records. The supported set of directives is defined relative to the CDNI Logging File Format version. The complete set of directives for version "cdni/1.0" are defined in Section 3.3. The directive list is not expected to require much extension, but when it does, the new directive MUST be defined and registered in the "CDNI Logging Directive Names" registry, as described in Figure 9, and a new version MUST be defined and registered in the "CDNI Logging File version" registry, as described in Section 6.2. For example, adding a new CDNI Logging Directive, e.g., "foo", to the set of directives defined for "cdni/1.0" in Section 3.3, would require registering both the new CDNI Logging Directive "foo" and a new CDNI Logging File version, e.g., "CDNI/2.0", which includes all of the existing CDNI Logging Directives of "cdni/1.0" plus "foo".

It is expected that as new logging requirements arise, the list of fields to log will change and expand. When adding new fields, the new fields MUST be defined and registered in the "CDNI Logging Field Names" registry, as described in Section 6.4, and a new record-type MUST be defined and registered in the "CDNI Logging record-types" registry, as described in Section 6.3. For example, adding a new CDNI Logging Field, e.g., "c-bar", to the set of fields defined for "cdni_http_request_v1" in Section 3.4.1, would require registering both the new CDNI Logging Field "c-bar" and a new CDNI record-type, e.g., "cdni_http_request_v2", which includes all of the existing CDNI Logging Fields of "cdni_http_request_v1" plus "c-bar".

3.6. CDNI Logging File Examples

Let us consider the upstream CDN and the downstream CDN labelled uCDN and dCDN-1 in Figure 1. When dCDN-1 acts as a downstream CDN for uCDN and performs content delivery on behalf of uCDN, dCDN-1 will include the CDNI Logging Records corresponding to the content deliveries performed on behalf of uCDN in the CDNI Logging Files for uCDN. An example CDNI Logging File communicated by dCDN-1 to uCDN is shown below in Figure 4.

```

#version:<HTAB>cdni/1.0<CRLF>

#UUID:<HTAB>urn:uuid:f81d4fae-7dec-11d0-a765-00a0c91e6bf6<CRLF>

#claimed-origin:<HTAB>cdni-logging-entity.dcdn-1.example.com<CRLF>

#record-type:<HTAB>cdni_http_request_v1<CRLF>

#fields:<HTAB>date<HTAB>time<HTAB>time-taken<HTAB>c-groupid<HTAB>
cs-method<HTAB>u-uri<HTAB>protocol<HTAB>
sc-status<HTAB>sc-total-bytes<HTAB>cs(User-Agent)<HTAB>
cs(Referer)<HTAB>s-cached<CRLF>

2013-05-17<HTAB>00:38:06.825<HTAB>9.058<HTAB>US/TN/MEM/38138<HTAB>
GET<HTAB>
http://cdni-ucdn.dcdn-1.example.com/video/movie100.mp4<HTAB>
HTTP/1.1<HTAB>200<HTAB>6729891<HTAB>"Mozilla/5.0
(Windows; U; Windows NT 6.0; en-US) AppleWebKit/533.4 (KHTML, like
Gecko) Chrome/5.0.375.127 Safari/533.4"<HTAB>
"host1.example.com"<HTAB>1<CRLF>

2013-05-17<HTAB>00:39:09.145<HTAB>15.32<HTAB>FR/PACA/NCE/06100<HTAB>
GET<HTAB>
http://cdni-ucdn.dcdn-1.example.com/video/movie118.mp4<HTAB>
HTTP/1.1<HTAB>200<HTAB>15799210<HTAB>"Mozilla/5.0
(Windows; U; Windows NT 6.0; en-US) AppleWebKit/533.4 (KHTML, like
Gecko) Chrome/5.0.375.127 Safari/533.4"<HTAB>
"host1.example.com"<HTAB>1<CRLF>

2013-05-17<HTAB>00:42:53.437<HTAB>52.879<HTAB>US/TN/MEM/38138<HTAB>
GET<HTAB>
http://cdni-ucdn.dcdn-1.example.com/video/picture11.mp4<HTAB>
HTTP/1.0<HTAB>200<HTAB>97234724<HTAB>"Mozilla/5.0
(Windows; U; Windows NT 6.0; en-US) AppleWebKit/533.4 (KHTML, like
Gecko) Chrome/5.0.375.127 Safari/533.4"<HTAB>
"host5.example.com"<HTAB>0<CRLF>

#SHA256-hash:<HTAB> 64-hexadecimal-digit hash value <CRLF>

```

Figure 4: CDNI Logging File Example

If uCDN establishes by some means (e.g., via TLS authentication when pulling the CDNI Logging File) the identity of the entity from which it pulled the CDNI Logging File, uCDN can add to the CDNI Logging an established-origin directive as illustrated below:

```

#established-origin:<HTAB>cdni-logging-entity.dcdn-
1.example.com<CRLF>

```


As illustrated in Figure 2, uCDN will then ingest the corresponding CDNI Logging Records into its Collection process, alongside the Logging Records generated locally by the uCDN itself. This allows uCDN to aggregate Logging Records for deliveries performed by itself (through Records generated locally) as well as for deliveries performed by its downstream CDN(s). This aggregate information can then be used (after Filtering and Rectification, as illustrated in Figure 2) by Log Consuming Applications that take into account deliveries performed by uCDN as well as by all of its downstream CDNs.

We observe that the time between

1. when a delivery is completed in dCDN and
2. when the corresponding Logging Record is ingested by the Collection process in uCDN

depends on a number of parameters such as the Logging Period agreed to by uCDN and dCDN, how much time uCDN waits before pulling the CDNI Logging File once it is advertised in the CDNI Logging Feed, and the time to complete the pull of the CDNI Logging File. Therefore, if we consider the set of Logging Records aggregated by the Collection process in uCDN in a given time interval, there could be a permanent significant timing difference between the CDNI Logging Records received from the dCDN and the Logging Records generated locally. For example, in a given time interval, the Collection process in uCDN may be aggregating Logging Records generated locally by uCDN for deliveries performed in the last hour and CDNI Logging Records generated in the dCDN for deliveries in the hour before last.

Say, that for some reason (for example a Surrogate bug), dCDN-1 could not collect the total number of bytes of the responses sent by the Surrogate (in other words, the value for sc-total-bytes is not available). Then the corresponding CDNI Logging records would contain a dash character ("-") in lieu of the value for the sc-total-bytes field (as specified in Section 3.4.1). In that case, the CDNI Logging File that would be communicated by dCDN-1 to uCDN is shown below in Figure 5.

```

#version:<HTAB>cdni/1.0<CRLF>

#UUID:<HTAB>urn:uuid:f81d4fae-7dec-11d0-a765-00a0c91e6bf6<CRLF>

#claimed-origin:<HTAB>cdni-logging-entity.dcdn-1.example.com<CRLF>

#record-type:<HTAB>cdni_http_request_v1<CRLF>

#fields:<HTAB>date<HTAB>time<HTAB>time-taken<HTAB>c-groupid<HTAB>
cs-method<HTAB>u-uri<HTAB>protocol<HTAB>
sc-status<HTAB>sc-total-bytes<HTAB>cs(User-Agent)<HTAB>
cs(Referer)<HTAB>s-cached<CRLF>

2013-05-17<HTAB>00:38:06.825<HTAB>9.058<HTAB>US/TN/MEM/38138<HTAB>
GET<HTAB>
http://cdni-ucdn.dcdn-1.example.com/video/movie100.mp4<HTAB>
HTTP/1.1<HTAB>200<HTAB>-<HTAB>"Mozilla/5.0
(Windows; U; Windows NT 6.0; en-US) AppleWebKit/533.4 (KHTML, like
Gecko) Chrome/5.0.375.127 Safari/533.4"<HTAB>
"host1.example.com"<HTAB>1<CRLF>

2013-05-17<HTAB>00:39:09.145<HTAB>15.32<HTAB>FR/PACA/NCE/06100<HTAB>
GET<HTAB>
http://cdni-ucdn.dcdn-1.example.com/video/movie118.mp4<HTAB>
HTTP/1.1<HTAB>200<HTAB>-<HTAB>"Mozilla/5.0
(Windows; U; Windows NT 6.0; en-US) AppleWebKit/533.4 (KHTML, like
Gecko) Chrome/5.0.375.127 Safari/533.4"<HTAB>
"host1.example.com"<HTAB>1<CRLF>

2013-05-17<HTAB>00:42:53.437<HTAB>52.879<HTAB>US/TN/MEM/38138<HTAB>
GET<HTAB>
http://cdni-ucdn.dcdn-1.example.com/video/picture11.mp4<HTAB>
HTTP/1.0<HTAB>200<HTAB>-<HTAB>"Mozilla/5.0
(Windows; U; Windows NT 6.0; en-US) AppleWebKit/533.4 (KHTML, like
Gecko) Chrome/5.0.375.127 Safari/533.4"<HTAB>
"host5.example.com"<HTAB>0<CRLF>

#SHA256-hash:<HTAB> 64-hexadecimal-digit hash value <CRLF>

```

Figure 5: CDNI Logging File Example With A Missing Field Value

3.7. Cascaded CDNI Logging Files Example

Let us consider the cascaded CDN scenario of uCDN, dCDN-2 and dCDN-3 as depicted in Figure 1. After completion of a delivery by dCDN-3 on behalf of dCDN-2, dCDN-3 will include a corresponding Logging Record in a CDNI Logging File that will be pulled by dCDN-2 and that is illustrated below in Figure 6. In practice, a CDNI Logging File is

likely to contain a very high number of CDNI Logging Records. However, for readability, the example in Figure 6 contains a single CDNI Logging Record.

```
#version:<HTAB>cdni/1.0<CRLF>

#UUID:<HTAB>urn:uuid:65718ef-0123-9876-adce4321bcde<CRLF>

#claimed-origin:<HTAB>cdni-logging-entity.dcdn-3.example.com<CRLF>

#record-type:<HTAB>cdni_http_request_v1<CRLF>

#fields:<HTAB>date<HTAB>time<HTAB>time-taken<HTAB>c-groupid<HTAB>
cs-method<HTAB>u-uri<HTAB>protocol<HTAB>
sc-status<HTAB>sc-total-bytes<HTAB>cs(User-Agent)<HTAB>
cs(Referer)<HTAB>s-cached<CRLF>

2013-05-17<HTAB>00:39:09.119<HTAB>14.07<HTAB>US/CA/SFO/94114<HTAB>
GET<HTAB>
http://cdni-dcdn-2.dcdn-3.example.com/video/movie118.mp4<HTAB>
HTTP/1.1<HTAB>200<HTAB>15799210<HTAB>"Mozilla/5.0
(Windows; U; Windows NT 6.0; en-US) AppleWebKit/533.4 (KHTML, like
Gecko) Chrome/5.0.375.127 Safari /533.4"<HTAB>
"host1.example.com"<HTAB>1<CRLF>

#SHA256-hash:<HTAB> 64-hexadecimal-digit hash value <CRLF>
```

Figure 6: Cascaded CDNI Logging File Example (dCDN-3 to dCDN-2)

If dCDN-2 establishes by some means (e.g., via TLS authentication when pulling the CDNI Logging File) the identity of the entity from which it pulled the CDNI Logging File, dCDN-2 can add to the CDNI Logging an established-origin directive as illustrated below:

```
#established-origin:<HTAB>cdni-logging-entity.dcdn-
3.example.com<CRLF>
```

dCDN-2 (behaving as an upstream CDN from the viewpoint of dCDN-3) will then ingest the CDNI Logging Record for the considered dCDN-3 delivery into its Collection process (as illustrated in Figure 2). This Logging Record may be aggregated with Logging Records generated locally by dCDN-2 for deliveries performed by dCDN-2 itself. Say, for illustration, that the content delivery performed by dCDN-3 on behalf of dCDN-2 had actually been redirected to dCDN-2 by uCDN, and say that another content delivery has just been redirected by uCDN to dCDN-2 and that dCDN-2 elected to perform the corresponding delivery itself. Then after Filtering and Rectification (as illustrated in Figure 2), dCDN-2 will include the two Logging Records corresponding

respectively to the delivery performed by dCDN-3 and the delivery performed by dCDN-2, in the next CDNI Logging File that will be communicated to uCDN. An example of such CDNI Logging File is illustrated below in Figure 7.

```
#version:<HTAB>cdni/1.0<CRLF>

#UUID:<HTAB>urn:uuid:1234567-8fedc-abab-0987654321ff<CRLF>

#claimed-origin:<HTAB>cdni-logging-entity.dcdn-2.example.com<CRLF>

#record-type:<HTAB>cdni_http_request_v1<CRLF>

#fields:<HTAB>date<HTAB>time<HTAB>time-taken<HTAB>c-groupid<HTAB>
cs-method<HTAB>u-uri<HTAB>protocol<HTAB>
sc-status<HTAB>sc-total-bytes<HTAB>cs(User-Agent)<HTAB>
cs(Referer)<HTAB>s-cached<CRLF>

2013-05-17<HTAB>00:39:09.119<HTAB>14.07<HTAB>US/CA/SFO/94114<HTAB>
GET<HTAB>
http://cdni-ucdn.dcdn-2.example.com/video/movie118.mp4<HTAB>
HTTP/1.1<HTAB>200<HTAB>15799210<HTAB>"Mozilla/5.0
(Windows; U; Windows NT 6.0; en-US) AppleWebKit/533.4 (KHTML, like
Gecko) Chrome/5.0.375.127 Safari /533.4"<HTAB>
"host1.example.com"<HTAB>1<CRLF>

2013-05-17<HTAB>01:42:53.437<HTAB>52.879<HTAB>FR/IDF/PAR/75001<HTAB>
GET<HTAB>
http://cdni-ucdn.dcdn-2.example.com/video/picture11.mp4<HTAB>
HTTP/1.0<HTAB>200<HTAB>97234724<HTAB>"Mozilla/5.0
(Windows; U; Windows NT 6.0; en-US) AppleWebKit/533.4 (KHTML, like
Gecko) Chrome/5.0.375.127 Safari /533.4"<HTAB>
"host5.example.com"<HTAB>0<CRLF>

#SHA256-hash:<HTAB> 64-hexadecimal-digit hash value <CRLF>
```

Figure 7: Cascaded CDNI Logging File Example (dCDN-2 to uCDN)

If uCDN establishes by some means (e.g., via TLS authentication when pulling the CDNI Logging File) the identity of the entity from which it pulled the CDNI Logging File, uCDN can add to the CDNI Logging an established-origin directive as illustrated below:

```
#established-origin:<HTAB>cdni-logging-entity.dcdn-
2.example.com<CRLF>
```

In the example of Figure 7, we observe that:

- o the first Logging Record corresponds to the Logging Record communicated earlier to dCDN-2 by dCDN-3, which corresponds to a delivery redirected by uCDN to dCDN-2 and then redirected by dCDN-2 to dCDN-3. The fields values in this Logging Record are copied from the corresponding CDNI Logging REcord communicated to dCDN2 by dCDN-3, with the exception of the u-uri that now reflects the URI convention between uCDN and dCDN-2 and that presents the delivery to uCDN as if it was performed by dCDN-2 itself. This reflects the fact that dCDN-2 had taken the full responsibility of the corresponding delivery (even if in this case, dCDN-2 elected to redirect the delivery to dCDN-3 so it is actually performed by dCDN-3 on behalf of dCDN-2).
- o the second Logging Record corresponds to a delivery redirected by uCDN to dCDN-2 and performed by dCDN-2 itself. The time of the delivery in this Logging Record may be significantly more recent than the first Logging Record since it was generated locally while the first Logging Record was generated by dCDN-3 and had to be advertised , and then pulled and then ingested into the dCDN-2 Collection process, before being aggregated with the second Logging Record.

4. Protocol for Exchange of CDNI Logging File After Full Collection

This section specifies a protocol for the exchange of CDNI Logging Files as specified in Section 3 after the CDNI Logging File is fully collected by the dCDN.

This protocol comprises:

- o a CDNI Logging feed, allowing the dCDN to notify the uCDN about the CDNI Logging Files that can be retrieved by that uCDN from the dCDN, as well as all the information necessary for retrieving each of these CDNI Logging Files. The CDNI Logging feed is specified in Section 4.1.
- o a CDNI Logging File pull mechanism, allowing the uCDN to obtain from the dCDN a given CDNI Logging File at the uCDN's convenience. The CDNI Logging File pull mechanisms is specified in Section 4.2.

An implementation of the CDNI Logging interface on the dCDN side (the entity generating the CDNI Logging file) MUST support the server side of the CDNI Logging feed (as specified in Section 4.1) and the server side of the CDNI Logging pull mechanism (as specified in Section 4.2).

An implementation of the CDNI Logging interface on the uCDN side (the entity consuming the CDNI Logging file) MUST support the client side

of the CDNI Logging feed (as specified in Section 4.1) and the client side of the CDNI Logging pull mechanism (as specified in Section 4.2).

4.1. CDNI Logging Feed

The server-side implementation of the CDNI Logging feed MUST produce an Atom feed [RFC4287]. This feed is used to advertise log files that are available for the client-side to retrieve using the CDNI Logging pull mechanism.

4.1.1. Atom Formatting

A CDNI Logging feed MUST be structured as an Archived feed, as defined in [RFC5005], and MUST be formatted in Atom [RFC4287]. This means it consists of a subscription document that is regularly updated as new CDNI Logging Files become available, and information about older CDNI Logging files is moved into archive documents. Once created, archive documents are never modified.

Each CDNI Logging File listed in an Atom feed MUST be described in an `atom:entry` container element.

The `atom:entry` MUST contain an `atom:content` element whose `"src"` attribute is a link to the CDNI Logging File and whose `"type"` attribute is the MIME Media Type indicating that the entry is a CDNI logging file. This MIME Media Type is defined as `"application/cdni"` (See [RFC7736]) with the Payload Type (`ptype`) parameter set to `"logging-file"`.

For compatibility with some Atom feed readers the `atom:entry` MAY also contain an `atom:link` entry whose `"href"` attribute is a link to the CDNI Logging File and whose `"type"` attribute is the MIME Media Type indicating that the entry is a CDNI Logging File using the `"application/cdni"` MIME Media Type with the Payload Type (`ptype`) parameter set to `"logging-file"` (See [RFC7736]).

The URI used in the `atom:id` of the `atom:entry` MUST contain the UUID of the CDNI Logging File.

The `atom:updated` in the `atom:entry` MUST indicate the time at which the CDNI Logging File was last updated.

4.1.2. Updates to Log Files and the Feed

CDNI Logging Files MUST NOT be modified by the dCDN once published in the CDNI Logging feed.

The frequency with which the subscription feed is updated, the period of time covered by each CDNI Logging File or each archive document, and timeliness of publishing of CDNI Logging Files are outside the scope of the present document and are expected to be agreed upon by uCDN and dCDN via other means (e.g., human agreement).

The server-side implementation **MUST** be able to set, and **SHOULD** set, HTTP cache control headers on the subscription feed to indicate the frequency at which the client-side is to poll for updates.

The client-side **MAY** use HTTP cache control headers (set by the server-side) on the subscription feed to determine the frequency at which to poll for updates. The client-side **MAY** instead, or in addition, use other information to determine when to poll for updates (e.g., a polling frequency that may have been negotiated between the uCDN and dCDN by mechanisms outside the scope of the present document and that is to override the indications provided in the HTTP cache control headers).

The potential retention limits (e.g., sliding time window) within which the dCDN is to retain and be ready to serve an archive document is outside the scope of the present document and is expected to be agreed upon by uCDN and dCDN via other means (e.g., human agreement). The server-side implementation **MUST** retain, and be ready to serve, any archive document within the agreed retention limits. Outside these agreed limits, the server-side implementation **MAY** indicate its inability to serve (e.g., with HTTP status code 404) an archive document or **MAY** refuse to serve it (e.g., with HTTP status code 403 or 410).

4.1.3. Redundant Feeds

The server-side implementation **MAY** present more than one CDNI Logging feed for redundancy. Each CDNI Logging File **MAY** be published in more than one feed.

A client-side implementation **MAY** support such redundant CDNI Logging feeds. If it supports redundant CDNI Logging feed, the client-side can use the UUID of the CDNI Logging File, presented in the atom:id element of the Atom feed, to avoid unnecessarily pulling and storing a given CDNI Logging File more than once.

4.1.4. Example CDNI Logging Feed

Figure 8 illustrates an example of the subscription document of a CDNI Logging feed.

```

<?xml version="1.0" encoding="utf-8"?>
<feed xmlns="http://www.w3.org/2005/Atom">
  <title type="text">CDNI Logging Feed</title>
  <updated>2013-03-23T14:46:11Z</updated>
  <id>urn:uuid:663ae677-40fb-e99a-049d-c5642916b8ce</id>
  <link href="https://dcdn.example/logfeeds/ucdn1"
    rel="self" type="application/atom+xml" />
  <link href="https://dcdn.example/logfeeds/ucdn1"
    rel="current" type="application/atom+xml" />
  <link href="https://dcdn.example/logfeeds/ucdn1/201303231400"
    rel="prev-archive" type="application/atom+xml" />
  <generator version="example version 1">CDNI Log Feed
    Generator</generator>
  <author><name>dcdn.example</name></author>
  <entry>
    <title type="text">CDNI Logging File for uCDN at
      2013-03-23 14:15:00</title>
    <id>urn:uuid:12345678-1234-abcd-00aa-01234567abcd</id>
    <updated>2013-03-23T14:15:00Z</updated>
    <content src="https://dcdn.example/logs/ucdn/
      http-requests-20130323141500000000"
      type="application/cdni"
      ptype="logging-file"/>
    <summary>CDNI Logging File for uCDN at
      2013-03-23 14:15:00</summary>
  </entry>
  <entry>
    <title type="text">CDNI Logging File for uCDN at
      2013-03-23 14:30:00</title>
    <id>urn:uuid:87654321-4321-dcba-aa00-dcba7654321</id>
    <updated>2013-03-23T14:30:00Z</updated>
    <content src="https://dcdn.example/logs/ucdn/
      http-requests-20130323143000000000"
      type="application/cdni"
      ptype="logging-file"/>
    <summary>CDNI Logging File for uCDN at
      2013-03-23 14:30:00</summary>
  </entry>
  ...
  <entry>
    ...
  </entry>
</feed>

```

Figure 8: Example subscription document of a CDNI Logging Feed

4.2. CDNI Logging File Pull

A client-side implementation of the CDNI Logging interface MAY pull, at its convenience, a CDNI Logging File that is published by the server-side in the CDNI Logging Feed (in the subscription document or an archive document). To do so, the client-side:

- o MUST implement HTTP/1.1 ([RFC7230],[RFC7231], [RFC7232], [RFC7233], [RFC7234], [RFC7235]), MAY also support other HTTP versions (e.g., HTTP/2 [RFC7540]) and MAY negotiate which HTTP version is actually used. This allows operators and implementers to choose to use later versions of HTTP to take advantage of new features, while still ensuring interoperability with systems that only support HTTP/1.1.
- o MUST use the URI that was associated to the CDNI Logging File (within the "src" attribute of the corresponding atom:content element) in the CDNI Logging Feed;
- o MUST support exchange of CDNI Logging Files with no content encoding applied to the representation;
- o MUST support exchange of CDNI Logging Files with "gzip" content encoding (as defined in [RFC7230]) applied to the representation.

Note that a client-side implementation of the CDNI Logging interface MAY pull a CDNI Logging File that it has already pulled.

The server-side implementation MUST respond to valid pull request by a client-side implementation for a CDNI Logging File published by the server-side in the CDNI Logging Feed (in the subscription document or an archive document). The server-side implementation:

- o MUST implement HTTP/1.1 to handle the client-side request and MAY also support other HTTP versions (e.g., HTTP/2);
- o MUST include the CDNI Logging File identified by the request URI inside the body of the HTTP response;
- o MUST support exchange of CDNI Logging Files with no content encoding applied to the representation;
- o MUST support exchange of CDNI Logging Files with "gzip" content encoding (as defined in [RFC7231]) applied to the representation.

Content negotiation approaches defined in [RFC7231] (e.g., using Accept-Encoding request-header field or Content-Encoding entity-header field) MAY be used by the client-side and server-side

implementations to establish the content-coding to be used for a particular exchange of a CDNI Logging File.

Applying compression content encoding (such as "gzip") is expected to mitigate the impact of exchanging the large volumes of logging information expected across CDNs. This is expected to be particularly useful in the presence of HTTP Adaptive Streaming (HAS) which, as per the present version of the document, will result in a separate CDNI Log Record for each HAS segment delivery in the CDNI Logging File.

The potential retention limits (e.g., sliding time window, maximum aggregate file storage quotas) within which the dCDN is to retain and be ready to serve a CDNI Logging File previously advertised in the CDNI Logging Feed is outside the scope of the present document and is expected to be agreed upon by uCDN and dCDN via other means (e.g., human agreement). The server-side implementation **MUST** retain, and be ready to serve, any CDNI Logging File within the agreed retention limits. Outside these agreed limits, the server-side implementation **MAY** indicate its inability to serve (e.g., with HTTP status code 404) a CDNI Logging File or **MAY** refuse to serve it (e.g., with HTTP status code 403 or 410).

5. Protocol for Exchange of CDNI Logging File During Collection

We note that, in addition to the CDNI Logging File exchange protocol specified in Section 4, implementations of the CDNI Logging interface may also support other mechanisms to exchange CDNI Logging Files. In particular, such mechanisms might allow the exchange of the CDNI Logging File to start before the file is fully collected. This can allow CDNI Logging Records to be communicated by the dCDN to the uCDN as they are gathered by the dCDN without having to wait until all the CDNI Logging Records of the same logging period are collected in the corresponding CDNI Logging File. This approach is commonly referred to as "tailing" of the file.

Such an approach could be used, for example, to exchange logging information with a significantly reduced time-lag (e.g., sub-minute or sub-second) between when the event occurred in the dCDN and when the corresponding CDNI Logging Record is made available to the uCDN. This can satisfy log-consuming applications requiring extremely fresh logging information such as near-real-time content delivery monitoring. Such mechanisms are for further study and outside the scope of this document.

6. IANA Considerations

6.1. CDNI Logging Directive Names Registry

The IANA is requested to create a new "CDNI Logging Directive Names" subregistry under the "Content Delivery Networks Interconnection (CDNI) Parameters" registry.

The initial contents of the CDNI Logging Directives registry comprise the names of the directives specified in Section 3.3 of the present document, and are as follows:

Directive Name	Reference
version	RFC xxxx
UUID	RFC xxxx
claimed-origin	RFC xxxx
established-origin	RFC xxxx
remark	RFC xxxx
record-type	RFC xxxx
fields	RFC xxxx
SHA256-hash	RFC xxxx

Figure 9

[Instructions to IANA: Replace "RFC xxxx" above by the RFC number of the present document]

Within the registry, names are to be allocated by IANA according to the "Specification Required" policy specified in [RFC5226]. Directive names are to be allocated by IANA with a format of NAMEFORMAT (see Section 3.1). All directive names defined in the logging file are case-insensitive as per the basic ABNF([RFC5234]).

Each specification that defines a new CDNI Logging directive needs to contain a description for the new directive with the same set of information as provided in Section 3.3 (i.e., format, directive value and occurrence).

6.2. CDNI Logging File version Registry

The IANA is requested to create a new "CDNI Logging File version" subregistry under the "Content Delivery Networks Interconnection (CDNI) Parameters" registry.

The initial contents of the CDNI Logging Logging File version registry comprise the value "cdni/1.0" specified in Section 3.3 of the present document, and are as follows:

version	Reference	Description
cdni/1.0	RFC xxxx	CDNI Logging File version 1.0 as specified in RFC xxxx

Figure 10

[Instructions to IANA: Replace "RFC xxxx" above by the RFC number of the present document]

Within the registry, version values are to be allocated by IANA according to the "Specification Required" policy specified in [RFC5226]. Version values are to be allocated by IANA with a format of NAMEFORMAT (see Section 3.1). All version values defined in the logging file are case-insensitive as per the basic ABNF([RFC5234]).

6.3. CDNI Logging record-types Registry

The IANA is requested to create a new "CDNI Logging record-types" subregistry under the "Content Delivery Networks Interconnection (CDNI) Parameters" registry.

The initial contents of the CDNI Logging record-types registry comprise the names of the CDNI Logging Record types specified in Section 3.4 of the present document, and are as follows:

record-types	Reference	Description
cdni_http_request_v1	RFC xxxx	CDNI Logging Record version 1 for content delivery using HTTP

Figure 11

[Instructions to IANA: Replace "RFC xxxx" above by the RFC number of the present document]

Within the registry, record-types are to be allocated by IANA according to the "Specification Required" policy specified in [RFC5226]. Record-types are to be allocated by IANA with a format of

NAMEFORMAT (see Section 3.1). All record-types defined in the logging file are case-insensitive as per the basic ABNF([RFC5234]).

Each specification that defines a new record-type needs to contain a description for the new record-type with the same set of information as provided in Section 3.4.1. This includes:

- o a list of all the CDNI Logging fields that can appear in a CDNI Logging Record of the new record-type
- o for all these fields: a specification of the occurrence for each Field in the new record-type
- o for every newly defined Field, i.e., for every Field that results in a registration in the CDNI Logging Field Names Registry (Section 6.4): a specification of the field name, format and field value.

6.4. CDNI Logging Field Names Registry

The IANA is requested to create a new "CDNI Logging Field Names" subregistry under the "Content Delivery Networks Interconnection (CDNI) Parameters" registry.

This registry is intended to be shared across the currently defined record-type (i.e., `cdni_http_request_v1`) as well as potential other CDNI Logging record-types that may be defined in separate specifications. When a Field from this registry is used by another CDNI Logging record-type, it is to be used with the exact semantics and format specified in the document that registered this field and that is identified in the Reference column of the registry. If another CDNI Logging record-type requires a Field with semantics that are not strictly identical, or a format that is not strictly identical then this new Field is to be registered in the registry with a different Field name. When a Field from this registry is used by another CDNI Logging record-type, it can be used with different occurrence rules.

The initial contents of the CDNI Logging fields Names registry comprise the names of the CDNI Logging fields specified in Section 3.4 of the present document, and are as follows:

Field Name	Reference
date	RFC xxxx
time	RFC xxxx
time-taken	RFC xxxx
c-groupid	RFC xxxx
s-ip	RFC xxxx
s-hostname	RFC xxxx
s-port	RFC xxxx
cs-method	RFC xxxx
cs-uri	RFC xxxx
u-uri	RFC xxxx
protocol	RFC xxxx
sc-status	RFC xxxx
sc-total-bytes	RFC xxxx
sc-entity-bytes	RFC xxxx
cs(insert_HTTP_header_name_here)	RFC xxxx
sc(insert_HTTP_header_name_here)	RFC xxxx
s-ccid	RFC xxxx
s-sid	RFC xxxx
s-cached	RFC xxxx

Figure 12

[Instructions to IANA: Replace "RFC xxxx" above by the RFC number of the present document]

Within the registry, names are to be allocated by IANA according to the "Specification Required" policy specified in [RFC5226]. Field names are to be allocated by IANA with a format of NHTABSTRING (see Section 3.1). All field names defined in the logging file are case-insensitive as per the basic ABNF([RFC5234]).

6.5. CDNI Logging MIME Media Type

The IANA is requested to register the following new Payload Type in the CDNI Payload Type registry for use with the application/cdni MIME media type.

[RFC Editor Note: Please replace the references to [RFCthis] below with this document's RFC number before publication.]

Payload Type	Specification
logging-file	[RFCthis]

Figure 13: MIME Media Type payload

The purpose of the logging-file payload type is to distinguish between CDNI Logging Files and other CDNI messages.

Interface: LI

Encoding: see Section 3.2, Section 3.3 and Section 3.4

7. Security Considerations

7.1. Authentication, Authorization, Confidentiality, Integrity Protection

An implementation of the CDNI Logging interface MUST support TLS transport of the CDNI Logging feed (Section 4.1) and of the CDNI Logging File pull (Section 4.2) as per [RFC2818] and [RFC7230].

TLS MUST be used by the server-side and the client-side of the CDNI Logging feed, as well as the server-side and the client-side of the CDNI Logging File pull mechanism, including authentication of the remote end, unless alternate methods are used for ensuring the security of the information exchanged over the LI interface (such as setting up an IPsec tunnel between the two CDNs or using a physically secured internal network between two CDNs that are owned by the same corporate entity).

The use of TLS for transport of the CDNI Logging feed and CDNI Logging File pull allows:

- o the dCDN and uCDN to authenticate each other using TLS client auth and TLS server auth.

and, once they have mutually authenticated each other, it allows:

- o the dCDN and uCDN to authorize each other (to ensure they are transmitting/receiving CDNI Logging File to/from an authorized CDN)
- o the CDNI Logging information to be transmitted with confidentiality

- o the integrity of the CDNI Logging information to be protected during the exchange.

When TLS is used, the general TLS usage guidance in [RFC7525] MUST be followed.

The SHA256-hash directive inside the CDNI Logging File provides additional integrity protection, this time targeting potential corruption of the CDNI logging information during the CDNI Logging File generation, storage or exchange. This mechanism does not itself allow restoration of the corrupted CDNI Logging information, but it allows detection of such corruption and therefore triggering of appropriate corrective actions (e.g., discard of corrupted information, attempt to re-obtain the CDNI Logging information). Note that the SHA256-hash does not protect against tampering by a third party, since such a third party could have recomputed and updated the SHA256-hash after tampering. Protection against third party tampering, when the CDNI Logging File is communicated over the CDN Logging Interface, can be achieved as discussed above through the use of TLS.

7.2. Denial of Service

This document does not define specific mechanism to protect against Denial of Service (DoS) attacks on the Logging Interface. However, the CDNI Logging feed and CDNI Logging pull endpoints are typically to be accessed only by a very small number of valid remote endpoints and therefore can be easily protected against DoS attacks through the usual conventional DOS protection mechanisms such as firewalling or use of Virtual Private Networks (VPNs).

Protection of dCDN Surrogates against spoofed delivery requests is outside the scope of the CDNI Logging interface.

7.3. Privacy

CDNs have the opportunity to collect detailed information about the downloads performed by End Users. A dCDN is expected to collect such information into CDNI Logging Files, which are then communicated to an uCDN.

Having detailed CDNI logging information known by the dCDN in itself does not represent a particular privacy concern since the dCDN is obviously fully aware of all information logged since it generated the information in the first place.

Transporting detailed CDNI logging information over the HTTP based CDNI Logging Interface does not represent a particular privacy

concern because it is protected by usual IETF privacy-protection mechanism (e.g., TLS).

When HTTP redirection is used between the uCDN and the dCDN, making detailed CDNI logging information known to the uCDN does not represent a particular privacy concern because the uCDN is already exposed at request redirection time to most of the information that shows up as CDNI logging information (e.g., enduser IP@, URL, HTTP headers). When DNS redirection is used between the uCDN and the dCDN, there are cases where there is no privacy concern in making detailed CDNI logging information known to the uCDN; this may be the case, for example, where (1) it is considered that because the uCDN has the authority (with respect to the CSP) and control on how the requests are delivered (including whether it is served by the uCDN itself or by a dCDN), the uCDN is entitled to access all detailed information related to the corresponding deliveries, and (2) there is no legal reasons to restrict access by the uCDN to all these detailed information. Conversely, still when DNS redirection is used between the uCDN and the dCDN, there are cases where there may be some privacy concern in making detailed CDNI logging information known to the uCDN; this may be the case, for example, because the uCDN is in a different jurisdiction to the dCDN resulting in some legal reasons to restrict access by the uCDN to all the detailed information related to the deliveries. In this latter case, the privacy concern can be taken into account when the uCDN and dCDN agree about which fields are to be conveyed inside the CDNI Logging Files and which privacy protection mechanism is to be used as discussed in the definition of the c-groupid field specified in Section 3.4.1.

Another privacy concern arises from the fact that large volumes of detailed information about content delivery to users, potentially traceable back to individual users, may be collected in CDNI Logging files. These CDNI Logging files represent high-value targets, likely concentrated in a fairly centralised system (although the CDNI Logging architecture does not mandate a particular level of centralisation/distribution) and at risk of potential data exfiltration. Note that the means of such data exfiltration are beyond the scope of the CDNI Logging interface itself (e.g., corrupted employee, corrupted logging storage system,...). This privacy concern calls for some protection.

The collection of large volumes of such information into CDNI Logging Files introduces potential End Users privacy protection concerns. Mechanisms to address these concerns are discussed in the definition of the c-groupid field specified in Section 3.4.1.

The use of mutually authenticated TLS to establish a secure session for the transport of the CDNI Logging feed and CDNI Logging pull as

discussed in Section 7.1 provides confidentiality while the logging information is in transit and prevents any party other than the authorised uCDN to gain access to the logging information.

We also note that the query string portion of the URL that may be conveyed inside the cs-uri and u-uri fields of CDNI Logging Files, or the HTTP cookies([RFC6265]) that may be conveyed as part of the cs(<HTTP-header-name>) field of CDNI Logging files, may contain personal information or information that can be exploited to derive personal information. Where this is a concern, the CDNI Logging interface specification allows the dCDN to not include the cs-uri and to include a u-uri that removes (or hides) the sensitive part of the query string and allows the dCDN to not include the cs(<HTTP-header-name>) fields corresponding to HTTP headers associated with cookies.

8. Acknowledgments

This document borrows from the W3C Extended Log Format [ELF].

Rob Murray significantly contributed into the text of Section 4.1.

The authors thank Ben Niven-Jenkins, Kevin Ma, David Mandelberg and Ray van Brandenburg for their ongoing input.

Brian Trammel and Rich Salz made significant contributions into making this interface privacy-friendly.

Finally, we also thank Sebastien Cubaud, Pawel Grochocki, Christian Jacquenet, Yannick Le Louedec, Anne Marrec, Emile Stephan, Fabio Costa, Sara Oueslati, Yvan Massot, Renaud Edel, Joel Favier and the contributors of the EU FP7 OCEAN project for their input in the early versions of this document.

9. References

9.1. Normative References

- [AES] NIST, "Advanced Encryption Standard (AES)", August 2015, <FIPS 197>.
- [GCM] NIST, "Recommendation for Block Cipher Modes of Operation: Galois/Counter Mode (GCM) and GMAC", November 2007, <SP 800-38D>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.

- [RFC3339] Klyne, G. and C. Newman, "Date and Time on the Internet: Timestamps", RFC 3339, DOI 10.17487/RFC3339, July 2002, <<http://www.rfc-editor.org/info/rfc3339>>.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, DOI 10.17487/RFC3986, January 2005, <<http://www.rfc-editor.org/info/rfc3986>>.
- [RFC4122] Leach, P., Mealling, M., and R. Salz, "A Universally Unique IDentifier (UUID) URN Namespace", RFC 4122, DOI 10.17487/RFC4122, July 2005, <<http://www.rfc-editor.org/info/rfc4122>>.
- [RFC4287] Nottingham, M., Ed. and R. Sayre, Ed., "The Atom Syndication Format", RFC 4287, DOI 10.17487/RFC4287, December 2005, <<http://www.rfc-editor.org/info/rfc4287>>.
- [RFC4648] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", RFC 4648, DOI 10.17487/RFC4648, October 2006, <<http://www.rfc-editor.org/info/rfc4648>>.
- [RFC5005] Nottingham, M., "Feed Paging and Archiving", RFC 5005, DOI 10.17487/RFC5005, September 2007, <<http://www.rfc-editor.org/info/rfc5005>>.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, DOI 10.17487/RFC5226, May 2008, <<http://www.rfc-editor.org/info/rfc5226>>.
- [RFC5234] Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, DOI 10.17487/RFC5234, January 2008, <<http://www.rfc-editor.org/info/rfc5234>>.
- [RFC7230] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", RFC 7230, DOI 10.17487/RFC7230, June 2014, <<http://www.rfc-editor.org/info/rfc7230>>.
- [RFC7231] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content", RFC 7231, DOI 10.17487/RFC7231, June 2014, <<http://www.rfc-editor.org/info/rfc7231>>.

- [RFC7232] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Conditional Requests", RFC 7232, DOI 10.17487/RFC7232, June 2014, <<http://www.rfc-editor.org/info/rfc7232>>.
- [RFC7233] Fielding, R., Ed., Lafon, Y., Ed., and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Range Requests", RFC 7233, DOI 10.17487/RFC7233, June 2014, <<http://www.rfc-editor.org/info/rfc7233>>.
- [RFC7234] Fielding, R., Ed., Nottingham, M., Ed., and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Caching", RFC 7234, DOI 10.17487/RFC7234, June 2014, <<http://www.rfc-editor.org/info/rfc7234>>.
- [RFC7235] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Authentication", RFC 7235, DOI 10.17487/RFC7235, June 2014, <<http://www.rfc-editor.org/info/rfc7235>>.
- [RFC7525] Sheffer, Y., Holz, R., and P. Saint-Andre, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", BCP 195, RFC 7525, DOI 10.17487/RFC7525, May 2015, <<http://www.rfc-editor.org/info/rfc7525>>.
- [RFC7540] Belshe, M., Peon, R., and M. Thomson, Ed., "Hypertext Transfer Protocol Version 2 (HTTP/2)", RFC 7540, DOI 10.17487/RFC7540, May 2015, <<http://www.rfc-editor.org/info/rfc7540>>.
- [SHA-3] NIST, "SHA-3 STANDARD: PERMUTATION-BASED HASH AND EXTENDABLE OUTPUT FUNCTIONS", November 2001, <<http://dx.doi.org/10.6028/NIST.FIPS.202>>.

9.2. Informative References

- [CHAR_SET] "IANA Character Sets registry", <<http://www.iana.org/assignments/character-sets/character-sets.xml>>.
- [ELF] Phillip M. Hallam-Baker, and Brian Behlendorf, "Extended Log File Format, W3C (work in progress), WD-logfile-960323", <<http://www.w3.org/TR/WD-logfile.html>>.

- [I-D.ietf-cdni-metadata]
Niven-Jenkins, B., Murray, R., Caulfield, M., and K. Ma,
"CDN Interconnection Metadata", draft-ietf-cdni-
metadata-17 (work in progress), May 2016.
- [I-D.ietf-tls-rfc5246-bis]
Dierks, T. and E. Rescorla, "The Transport Layer Security
(TLS) Protocol Version 1.3", draft-ietf-tls-rfc5246-bis-00
(work in progress), April 2014.
- [I-D.snell-atompub-link-extensions]
Snell, J., "Atom Link Extensions", draft-snell-atompub-
link-extensions-09 (work in progress), June 2012.
- [RFC1945] Berners-Lee, T., Fielding, R., and H. Frystyk, "Hypertext
Transfer Protocol -- HTTP/1.0", RFC 1945,
DOI 10.17487/RFC1945, May 1996,
<<http://www.rfc-editor.org/info/rfc1945>>.
- [RFC2818] Rescorla, E., "HTTP Over TLS", RFC 2818,
DOI 10.17487/RFC2818, May 2000,
<<http://www.rfc-editor.org/info/rfc2818>>.
- [RFC5869] Krawczyk, H. and P. Eronen, "HMAC-based Extract-and-Expand
Key Derivation Function (HKDF)", RFC 5869,
DOI 10.17487/RFC5869, May 2010,
<<http://www.rfc-editor.org/info/rfc5869>>.
- [RFC6234] Eastlake 3rd, D. and T. Hansen, "US Secure Hash Algorithms
(SHA and SHA-based HMAC and HKDF)", RFC 6234,
DOI 10.17487/RFC6234, May 2011,
<<http://www.rfc-editor.org/info/rfc6234>>.
- [RFC6265] Barth, A., "HTTP State Management Mechanism", RFC 6265,
DOI 10.17487/RFC6265, April 2011,
<<http://www.rfc-editor.org/info/rfc6265>>.
- [RFC6707] Niven-Jenkins, B., Le Faucheur, F., and N. Bitar, "Content
Distribution Network Interconnection (CDNI) Problem
Statement", RFC 6707, DOI 10.17487/RFC6707, September
2012, <<http://www.rfc-editor.org/info/rfc6707>>.
- [RFC6770] Bertrand, G., Ed., Stephan, E., Burbridge, T., Eardley,
P., Ma, K., and G. Watson, "Use Cases for Content Delivery
Network Interconnection", RFC 6770, DOI 10.17487/RFC6770,
November 2012, <<http://www.rfc-editor.org/info/rfc6770>>.

- [RFC6983] van Brandenburg, R., van Deventer, O., Le Faucheur, F., and K. Leung, "Models for HTTP-Adaptive-Streaming-Aware Content Distribution Network Interconnection (CDNI)", RFC 6983, DOI 10.17487/RFC6983, July 2013, <<http://www.rfc-editor.org/info/rfc6983>>.
- [RFC7336] Peterson, L., Davie, B., and R. van Brandenburg, Ed., "Framework for Content Distribution Network Interconnection (CDNI)", RFC 7336, DOI 10.17487/RFC7336, August 2014, <<http://www.rfc-editor.org/info/rfc7336>>.
- [RFC7337] Leung, K., Ed. and Y. Lee, Ed., "Content Distribution Network Interconnection (CDNI) Requirements", RFC 7337, DOI 10.17487/RFC7337, August 2014, <<http://www.rfc-editor.org/info/rfc7337>>.
- [RFC7736] Ma, K., "Content Delivery Network Interconnection (CDNI) Media Type Registration", RFC 7736, DOI 10.17487/RFC7736, December 2015, <<http://www.rfc-editor.org/info/rfc7736>>.

Authors' Addresses

Francois Le Faucheur (editor)
FR

Phone: +33 6 19 98 50 90
Email: flefauch@gmail.com

Gilles Bertrand (editor)

Phone: +41 76 675 91 44
Email: gilbertrand@gmail.com

Iuniana Oprescu (editor)
FR

Email: iuniana.oprescu@gmail.com

Roy Peterkofsky
Google Inc.
345 Spear St, 4th Floor
San Francisco CA 94105
USA

Email: peterkofsky@google.com

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 24, 2014

B. Niven-Jenkins
R. Murray
G. Watson
Velocix (Alcatel-Lucent)
M. Caulfield
K. Leung
Cisco Systems
K. Ma
Azuki Systems, Inc.
October 21, 2013

CDN Interconnect Metadata
draft-ietf-cdni-metadata-03

Abstract

The CDNI Metadata Interface enables interconnected CDNs to exchange content distribution metadata in order to enable content acquisition and delivery. The CDNI metadata associated with a piece of content provides a downstream CDN with sufficient information for the downstream CDN to service content requests on behalf of an upstream CDN. This document describes both the core set of CDNI metadata and the protocol for exchanging that metadata.

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 April 24, 2014.

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	3
1.1. Terminology	4
2. Design Principles	4
3. CDNI Metadata Data Model	5
3.1. HostIndex, HostMetadata & PathMetadata objects	6
3.2. Generic CDNI Metadata Object Properties	9
3.3. Metadata Inheritance and Override	10
3.4. Metadata Naming	10
4. Encoding-Independent CDNI Metadata Object Descriptions	11
4.1. CDNI Metadata Structural Object Descriptions	11
4.1.1. HostIndex	12
4.1.2. HostMatch	12
4.1.3. HostMetadata	13
4.1.4. PathMatch	13
4.1.5. PathMetadata	14
4.1.6. PatternMatch	14
4.1.7. GenericMetadata	15
4.2. CDNI Metadata Property Object Descriptions	16
4.2.1. Source Metadata	16
4.2.1.1. Source	16
4.2.2. LocationACL Metadata	17
4.2.2.1. LocationRule	17
4.2.2.2. Footprint	17
4.2.3. TimeWindowACL Metadata	18
4.2.3.1. TimeWindowRule	18
4.2.3.2. TimeWindow	19
4.2.4. ProtocolACL Metadata	19
4.2.4.1. ProtocolRule	19
4.2.5. Authorization Metadata	20
4.2.6. Auth	20
4.2.6.1. Credentials Auth Type	21

4.2.7.	Cache	21
4.2.8.	Grouping	22
4.3.	CDNI Metadata Simple Data Type Descriptions	22
4.3.1.	Link	22
4.3.2.	Protocol	23
4.3.3.	Endpoint	23
4.3.4.	URI	23
4.3.5.	Time	23
5.	CDNI Metadata Capabilities	24
5.1.	Protocol ACL Capabilities	24
5.2.	Authorization Metadata Capabilities	24
6.	CDNI Metadata interface	25
6.1.	Transport	25
6.2.	Retrieval of CDNI Metadata resources	26
6.3.	Bootstrapping	27
6.4.	Encoding	28
6.4.1.	MIME Media Types	28
6.4.2.	JSON Encoding of Objects	28
6.4.2.1.	JSON Example	29
6.5.	Extensibility	33
6.5.1.	Metadata Enforcement	33
6.5.2.	Metadata Override	33
6.6.	Versioning	34
7.	IANA Considerations	34
7.1.	GenericMetadata Type Registry	34
7.1.1.	GenericMetadata Sub-Registries	35
7.1.1.1.	Footprint Sub-Registry	35
7.1.1.2.	Protocol Sub-Registry	36
7.1.1.3.	Authentication Sub-Registry	37
8.	Security Considerations	37
9.	Acknowledgements	37
10.	References	38
10.1.	Normative References	38
10.2.	Informative References	38
	Appendix A. Relationship to the CDNI Requirements	39
	Authors' Addresses	39

1. Introduction

CDNI enables a downstream CDN to service content requests on behalf of an upstream CDN. The CDNI metadata associated with a piece of content (or with a set of contents) provides a downstream CDN with sufficient information for servicing content requests on behalf of an upstream CDN in accordance with the policies defined by the upstream CDN.

The CDNI Metadata Interface is introduced by [RFC6707] along with three other interfaces that may be used to compose a CDNI solution

(Control, Request Routing and Logging). [I-D.ietf-cdni-framework] expands on the information provided in [RFC6707] and describes each interface, and the relationships between them, in more detail. The requirements for the CDNI metadata interface are specified in [I-D.ietf-cdni-requirements].

This document focuses on the CDNI Metadata interface which enables a downstream CDN to obtain CDNI Metadata from an upstream CDN so that the downstream CDN can properly process and respond to:

- o Redirection Requests received over the CDNI Request Routing protocol.
- o Content Requests received directly from User Agents.

Specifically this document proposes:

- o A data structure for mapping content requests to CDNI Metadata properties (Section 3).
- o An initial set of CDNI Metadata properties (Section 4.2).
- o A RESTful web service for the transfer of CDNI Metadata (Section 6).

1.1. Terminology

This document reuses the terminology defined in [RFC6707].

Additionally, the following terms are used throughout this document and are defined as follows:

- o Object - a collection of properties
- o Property - a key and value pair where the key is a property name and the value is the property value or an object.

2. Design Principles

The proposed CDNI Metadata Interface was designed to achieve the following objectives:

1. Cacheability of CDNI metadata objects
2. Deterministic mapping from redirection and content requests to CDNI metadata properties

3. Support for DNS redirection as well as application-specific redirection (for example HTTP redirection)
4. Minimal duplication of CDNI metadata
5. Leverage existing protocols

Cacheability improves the latency of acquiring metadata while maintaining its freshness and therefore improves the latency of serving content requests. The CDNI Metadata Interface uses HTTP to achieve cacheability.

Deterministic mappings from content to metadata properties eliminates ambiguity and ensures that policies are applied consistently by all downstream CDNs.

Support for both HTTP and DNS redirection ensures that the CDNI Metadata Interface can be used for HTTP and DNS redirection and also meets the same design principles for both HTTP and DNS based redirection schemes.

Minimal duplication of CDNI metadata provides space efficiency on storage in the CDNs, on caches in the network, and across the network between CDNs.

Leveraging existing protocols avoids reinventing common mechanisms such as data structure encoding (e.g. XML, JSON) and data transport (e.g. HTTP).

3. CDNI Metadata Data Model

The CDNI Metadata Model describes a data structure for mapping redirection requests and content requests to metadata properties. Metadata properties describe how to acquire, authorize, and deliver content from a downstream CDN. The data model relies on the assumption that these metadata properties may be aggregated based on the hostname of the content and subsequently on the resource path of the content. The data model associates a set of CDNI Metadata properties with a Hostname to form a default set of metadata properties for content delivered for that Hostname. That default set of metadata properties can be overridden by properties that apply to specific paths within a URI.

Different Hostnames and URI paths will be associated with different sets of CDNI Metadata properties in order to describe the required behaviour when a dCDN surrogate is processing User Agent requests for content at that Hostname or URI path. As a result of this structure, significant commonality may exist between the CDNI Metadata

properties specified for different Hostnames, different URI paths within a Hostname and different URI paths on different Hostnames. For example the definition of which User Agent IP addresses should be treated as being grouped together into a single network or geographic location is likely to be common for a number of different Hostnames. Another example is that although a uCDN is likely to have several different policies configured to express geo-blocking rules, it is likely that a single geo-blocking policy would be applied to multiple Hostnames delivered through the CDN.

In order to enable the CDNI Metadata for a given Hostname or URI Path to be decomposed into sets of CDNI Metadata properties that can be reused by multiple Hostnames and URI Paths, the CDNI Metadata interface specified in this document splits the CDNI Metadata into a number of objects. Efficiency is improved by enabling a single CDNI Metadata object (that is shared across Hostname and/or URI paths) to be retrieved by a dCDN once, even if it is referenced by the CDNI Metadata of multiple Hostnames.

Section 3.1 introduces a high level description of the HostIndex, HostMetadata and PathMetadata objects and describes the relationships between those objects.

Section 3.2 introduces a high level description of the CDNI GenericMetadata object which represents the level at which CDNI Metadata override occurs between HostMetadata and PathMetadata objects.

Section 4 describes in detail the specific CDNI Metadata objects and properties which may be contained within a CDNI GenericMetadata object.

3.1. HostIndex, HostMetadata & PathMetadata objects

A HostIndex object contains a list of Hostnames (and/or IP addresses) for which content requests may be delegated to the downstream CDN. The HostIndex is the starting point for accessing the uCDN's CDNI Metadata data store. It enables surrogates in the dCDN to deterministically discover, on receipt of a User Agent request for content, which other CDNI Metadata objects it requires in order to deliver the requested content.

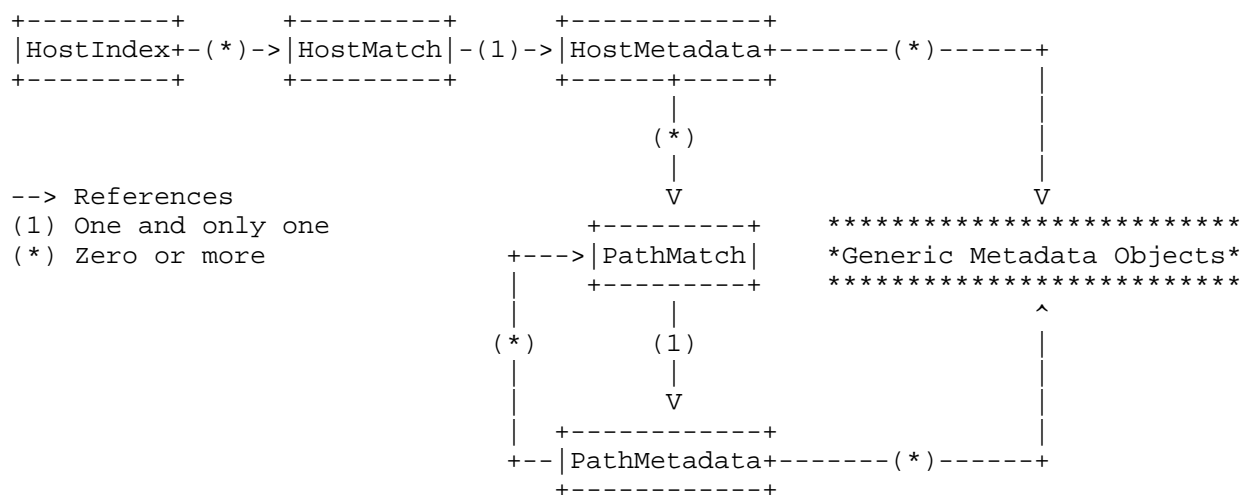
The HostIndex links Hostnames (and/or IP addresses) to HostMetadata objects via HostMatch objects. HostMetadata objects contain (or reference) the default CDNI Metadata required to serve content for that host. When looking up CDNI Metadata, the downstream CDN looks up the requested Hostname (or IP address) in the HostIndex, from there it can find HostMetadata which describes properties for a host

and PathMetadata which may override those properties for given URI paths within the host.

As well as containing the default CDNI Metadata for the specified Hostname, HostMetadata and PathMetadata objects may also contain PathMatch objects which in turn contain PathMetadata objects. PathMatch objects override the CDNI Metadata in the HostMetadata object or one or more preceding PathMetadata objects with more specific CDNI Metadata that applies to content requests matching the pattern defined in that PathMatch object.

For the purposes of retrieving CDNI Metadata all other required CDNI Metadata objects and their properties are discoverable from the appropriate HostMetadata, PathMatch and PathMetadata objects for the requested content.

The relationships between the HostIndex, HostMatch, HostMetadata, PathMatch and PathMetadata objects are described in Figure 1.



Key: ----> = References

Figure 1: Relationships between the HostIndex, HostMetadata & PathMetadata CDNI Metadata Objects

The relationships in Figure 1 are summarised in Table 1 below.

Data Object	Objects it References
HostIndex	0 or more HostMatch objects.

HostMatch	1 HostMetadata object.
HostMetadata	0 or more PathMatch objects. 0 or more GenericMetadata objects.
PathMatch	1 PathMetadata object.
PathMetadata	0 or more PathMatch objects. 0 or more GenericMetadata objects.

Table 1: Relationships between CDNI Metadata Objects

The table below describes the HostIndex, HostMetadata and PathMetadata objects in more detail.

Data Object	Description
HostIndex	A HostIndex object lists HostMatch objects
HostMatch	A HostMatch object defines a hostname to match against a requested host, and contains or references a HostMetadata object which contains CDNI Metadata objects to be applied when a request matches against the hostname. For example, if "example.com" is a content provider, a HostMatch object may include an entry for "example.com" with the URI of the associated HostMetadata object.
HostMetadata	A HostMetadata object contains (or references) the default CDNI Metadata objects for content served from that host, i.e. the CDNI Metadata objects for content requests that do not match any of the PathMatch objects contained or referenced by that HostMetadata object. For example, a HostMetadata object may describe the metadata properties which apply to "example.com" and may contain PathMatches for "example.com/movies/*" and "example.com/music/*" which reference corresponding PathMetadata objects that contain the CDNI Metadata objects for those more specific URI paths.
PathMatch	A PathMatch object defines a pattern to match against the requested URI path, and contains or references a PathMetadata object which contains (or references) the CDNI Metadata objects to be applied when a content request matches against the defined URI path pattern.
PathMetadata	A PathMetadata object contains the CDNI GenericMetadata objects for content served with

	the associated URI path (defined in a PathMatch object). A PathMetadata object may also contain PathMatch objects in order to recursively define more specific URI paths that require different (e.g. more specific) CDNI Metadata to this one. For example, the PathMetadata object which applies to "example.com/movies/*" may describe CDNI Metadata which apply to that resource path and may contain a PathMatch object for "example.com/movies/hd/*" which would reference the corresponding PathMetadata object for the "example.com/movies/hd/" path prefix.
GenericMetadata	A GenericMetadata object contains individual CDNI Metadata objects which define the specific policies and attributes needed to properly deliver the associated content.

Table 2: HostIndex, HostMetadata and PathMetadata CDNI Metadata Objects

3.2. Generic CDNI Metadata Object Properties

The HostMetadata and PathMetadata objects contain or can reference other CDNI Metadata objects that contain properties which describe how User Agent requests for content should be processed, for example where to acquire the content, authorization rules that should be applied, delivery location restrictions and so on. Each such CDNI Metadata object is a specialization of a CDNI GenericMetadata object. The GenericMetadata object abstracts the basic information required for Metadata override and opaque Metadata distribution, from the specifics of any given property (e.g., property semantics, enforcement options, etc.).

The GenericMetadata object defines the type of properties contained within it as well as whether or not the properties are mandatory to enforce. If the dCDN does not understand or support the property type and the property type is mandatory to enforce, the dCDN MUST NOT serve the content to the User Agent. If the dCDN does not understand or support the property type it is also not going to be able to properly propagate the Metadata for cascaded distribution. If the dCDN does not understand or support the property type and the property type is not mandatory to enforce, then the GenericMetadata object may be safely ignored.

Although a CDN cannot serve content to a User Agent if a mandatory property cannot be enforced, it may be safe to redistribute that

metadata to another CDN without modification. For example, in the cascaded CDN case, a transit CDN may pass through mandatory-to-enforce metadata to the delivery CDN. For Metadata which does not require customization, the data representation received off the wire MAY be stored and redistributed without being natively understood or supported by the transit CDN. However, for Metadata which require translations, transparent redistribution of the uCDN Metadata values may not be appropriate. Certain Metadata may be safely, though possibly not optimally, redistributed unmodified, e.g., source acquisition address may not be optimal if transparently redistributed, but may still work. Redistribution safety MUST be specified for each GenericMetadata.

3.3. Metadata Inheritance and Override

In the data model, a HostMetadata object may contain (or reference) multiple PathMetadata objects (via PathMatch objects). Each PathMetadata object may in turn contain (or reference) other PathMetadata objects. HostMetadata and PathMetadata objects form an inheritance tree where each node in the tree inherits or overrides the property values set by its parent.

GenericMetadata objects of a given type override all GenericMetadata objects of the same type previously defined by any parent object in the tree. GenericMetadata objects of a given type previously defined by a parent object in the tree are inherited when no object of the same type is defined by the child object. For example, if HostMetadata for the host "example.com" contains GenericMetadata objects of type LocationACL and TimeWindowACL, while a PathMetadata object which applies to "example.com/movies/*" defines an alternate GenericMetadata object of type TimeWindowACL, then:

the TimeWindowACL defined in the PathMetadata would override the TimeWindowACL defined in the HostMetadata

the LocationACL defined in the HostMetadata would be inherited for all User Agent requests for content under "example.com/movies".

The PathMetadata defined TimeWindowACL would override the TimeWindowACL defined in the HostMetadata for all User Agent requests for movies.

3.4. Metadata Naming

GenericMetadata objects are identified by their type. The type SHOULD be descriptive, and MAY be hierarchical to support aggregating groups of properties for the purpose of readability and for avoiding name conflicts between vendor extensions. A dotted alpha-numeric notation is suggested for human readability.

Metadata types defined by this document are not hierarchical.

Examples of GenericMetadata object type names:

LocationACL

ext.vendor1.featurex

ext.vendor1.featurey

ext.vendor2.featurex

4. Encoding-Independent CDNI Metadata Object Descriptions

Section 4.1 provides the definitions of each object type declared in Section 3. These objects are described as structural objects as they provide the structure for the inheritance tree and identifying which specific properties apply to a given User Agent content request.

Section 4.2 provides the definitions for the set of core metadata objects which may be contained within a GenericMetadata object. These objects are described as property objects as they define the semantics, enforcement options, and serialization rules for specific properties. These properties govern how User Agent requests for content are handled. Property objects may be composed of or contain references to other objects. In those cases the value of the property can be either an object of that type (the object is embedded) or a Link object that contains a URI and relationship that can be dereferenced to retrieve the CDNI Metadata object that represents the value of that property.

Note: In the following sections, the term "mandatory-to-specify" is used to convey which objects or properties must be specified for a given parent object or property. When mandatory-to-specify is set to true, it implies that if the parent object is specified, then the defined object or property MUST also be specified, e.g., a HostMatch object without a host to match against does not make sense, therefore, the host is mandatory-to-specify inside a parent HostMatch object.

4.1. CDNI Metadata Structural Object Descriptions

Each of the sub-sections below describe the structural objects defined in Table 2.

4.1.1. HostIndex

The HostIndex object is the entry point into the CDNI Metadata hierarchy. It contains a list of HostMatch objects. An incoming content request is matched against the hostname inside of each of the listed HostMatch objects to find the HostMatch object which applies to the request.

Property: hosts

Description: List of HostMatch objects, in priority order.

Type: List of HostMatch objects

Mandatory-to-Specify: Yes.

4.1.2. HostMatch

The HostMatch object contains a hostname or IP address to match against content requests. The HostMatch object also contains a reference to Metadata objects to apply if a match is found.

Property: host

Description: String (hostname or IP address) to match against the requested host.

Type: String

Mandatory-to-Specify: Yes.

Property: host-metadata

Description: CDNI Metadata to apply when delivering content that matches this host.

Type: HostMetadata

Mandatory-to-Specify: Yes.

4.1.3. HostMetadata

The HostMetadata object contains both Metadata that applies to content requests for a particular host and a list of pattern matches for finding more specific Metadata based on the resource path in a content request.

Property: metadata

Description: List of host related metadata.

Type: List of GenericMetadata objects

Mandatory-to-Specify: Yes.

Property: paths

Description: Path specific rules. First match applies.

Type: List of PathMatch objects

Mandatory-to-Specify: No.

4.1.4. PathMatch

The PathMatch object contains an expression given as a PatternMatch object to match against a resource URI path and Metadata objects to apply if a match is found.

Property: path-pattern

Description: Pattern to match against the requested path, i.e. against the [RFC3986] path-absolute.

Type: PatternMatch

Mandatory-to-Specify: Yes.

Property: path-metadata

Description: CDNI Metadata to apply when delivering content that matches this pattern.

Type: PathMetadata

Mandatory-to-Specify: Yes.

4.1.5. PathMetadata

A PathMetadata object contains the CDNI Metadata properties for content served with the associated URI path (defined in a PathMatch object). Note that if CDNI metadata is used as an input to CDNI request routing and DNS-based redirection is employed, then any metadata at the PathMetadata level or below will be inaccessible at request routing time.

Property: metadata

Description: List of path related metadata.

Type: List of GenericMetadata objects

Mandatory-to-Specify: Yes.

Property: paths

Description: Path specific rules. First match applies.

Type: List of PathMatch objects

Mandatory-to-Specify: No.

4.1.6. PatternMatch

A PatternMatch object contains the pattern string and flags that describe the PathMatch expression.

Property: pattern

Description: A pattern for string matching. The pattern may contain the wildcards * and ?, where * matches any sequence of characters (including the empty string) and ? matches exactly one character. The three literals \ , * and ? should be escaped as \\, * and \?

Type: String

Mandatory-to-Specify: Yes.

Property: case-sensitive

Description: Flag indicating whether or not case-sensitive matching should be used.

Type: Boolean

Mandatory-to-Specify: No. Default is case-insensitive match.

Property: ignore-query-string

Description: List of query parameters which should be ignored when searching for a pattern match. If all query parameters should be ignored then the list MUST be empty.

Type: List of String

Mandatory-to-Specify: No. Default is to include query strings when matching.

4.1.7. GenericMetadata

A GenericMetadata object is a abstraction for managing individual CDNI Metadata properties in an opaque manner.

Property: type

Description: CDNI Metadata property object type.

Type: String

Mandatory-to-Specify: Yes.

Property: value

Description: CDNI Metadata property object.

Type: matches the type property above

Mandatory-to-Specify: Yes.

Property: mandatory-to-enforce

Description: Flag identifying whether or not the enforcement of the property Metadata is required.

Type: Boolean

Mandatory-to-Specify: Yes.

Property: safe-to-redistribute

Description: Flag identifying whether or not the property Metadata may be safely redistributed without modification.

Type: Boolean

Mandatory-to-Specify: No. Default is allow transparent redistribution.

4.2. CDNI Metadata Property Object Descriptions

4.2.1. Source Metadata

Source Metadata provides the dCDN information about content acquisition e.g. how to contact an uCDN Surrogate or an Origin Server to obtain the content to be served. The sources are not necessarily the actual Origin Servers operated by the CSP but might be a set of Surrogates in the uCDN.

Property: sources

Description: Sources from which the dCDN can acquire content, listed in priority order.

Type: List of Source objects

Mandatory-to-Specify: No. Default is to use static configuration, out of band of the metadata interface.

4.2.1.1. Source

A Source object describes the Source which should be used by the dCDN for content acquisition, e.g. a Surrogate within the uCDN or an alternate Origin Server, the protocol to be used and any authentication method.

Property: auth

Description: Authentication method to use when requesting content from this source.

Type: Auth

Mandatory-to-Specify: No. Default is no authentication is required.

Property: endpoints

Description: Origins from which the dCDN can acquire content.

Type: List of EndPoint objects

Mandatory-to-Specify: Yes.

Property: protocol

Description: Network retrieval protocol to use when requesting content from this source.

Type: Protocol

Mandatory-to-Specify: Yes.

4.2.2. LocationACL Metadata

LocationACL Metadata defines location-based restrictions.

Property: locations

Description: Access control list which applies restrictions to delivery based on client location.

Type: List of LocationRule objects

Mandatory-to-Specify: No. Default is allow all locations.

4.2.2.1. LocationRule

A LocationRule contains or references a list of Location objects and the corresponding action.

Property: footprints

Description: List of footprints to which the rule applies.

Type: List of Footprint objects

Mandatory-to-Specify: Yes.

Property: action

Description: Defines whether the rule specifies locations to allow or deny.

Type: Enumeration [allow|deny]

Mandatory-to-Specify: No. Default is deny.

4.2.2.2. Footprint

A Footprint object describes the footprint to which a LocationRule may be applied by, e.g. an IPv4 address range or a geographic location.

Property: type

Description: Registered footprint type (see Section 7.1.1.1).

Type: String

Mandatory-to-Specify: Yes.

Property: value

Description: Footprint object conforming to the specification associated with the registered footprint type.

Type: String

Mandatory-to-Specify: Yes.

4.2.3. TimeWindowACL Metadata

TimeWindowACL Metadata defines time-based restrictions.

Property: times

Description: Access control list which applies restrictions to delivery based on request time.

Type: List of TimeWindowRule objects

Mandatory-to-Specify: No. Default is allow all time windows.

4.2.3.1. TimeWindowRule

A TimeWindowRule contains or references a list of TimeWindow objects and the corresponding action.

Property: times

Description: List of time windows to which the rule applies.

Type: List of TimeWindow objects

Mandatory-to-Specify: Yes.

Property: action

Description: Defines whether the rule specifies time windows to allow or deny.

Type: Enumeration [allow|deny]

Mandatory-to-Specify: No. Default is deny.

4.2.3.2. TimeWindow

A TimeWindow object describes a time range which may be applied by an ACLRule, e.g. Start 09:00AM 01/01/2000 UTC End 17:00PM 01/01/2000 UTC.

Property: start

Description: The start time of the window.

Type: Time

Mandatory-to-Specify: Yes.

Property: end

Description: The end time of the window.

Type: Time

Mandatory-to-Specify: Yes.

4.2.4. ProtocolACL Metadata

ProtocolACL Metadata defines delivery protocol restrictions.

Property: protocols

Description: Access control list which applies restrictions to delivery based on delivery protocol.

Type: List of ProtocolRule objects

Mandatory-to-Specify: No. Default is allow all protocols.

4.2.4.1. ProtocolRule

A ProtocolRule contains or references a list of Protocol objects. ProtocolRule objects are used to construct a ProtocolACL to apply restrictions to content acquisition or delivery.

Property: protocols

Description: List of protocols to which the rule applies.

Type: List of protocol objects

Mandatory-to-Specify: Yes.

Property: action

Description: Defines whether the rule specifies protocols to allow or deny.

Type: Enumeration [allow|deny]+

Mandatory-to-Specify: No. Default is allow all protocols.

Property: direction

Description: Defines whether the ProtocolRule specifies protocols for acquisition or delivery.

Type: Enumeration [acquisition|delivery]

Mandatory-to-Specify: No. Default is to apply the rule to both acquisition and delivery.

4.2.5. Authorization Metadata

Authorization Metadata define content authorization methods.

Property: methods

Description: Options for authenticating content requests. All options in the list are equally valid.

Type: List of Auth objects

Mandatory-to-Specify: No. Default is no authorization required.

4.2.6. Auth

An Auth object defines authentication and authorization methods to be used during content delivery and content acquisition.

Property: type

Description: Registered Auth type (see Section 7.1.1.3).

Type: String

Mandatory-to-Specify: Yes.

Property: value

Description: Auth object conforming to the specification associated with the registered Auth type.

Type: String

Mandatory-to-Specify: Yes.

4.2.6.1. Credentials Auth Type

Credentials Auth is a type of Auth object with type "credentials" (see Section 7.1.1.3). The CredentialsAuth object contains the following properties:

Property: username

Description: Identification of user.

Type: String

Mandatory-to-Specify: Yes.

Property: password

Description: Password for user identified by username property.

Type: String

Mandatory-to-Specify: Yes.

4.2.7. Cache

A Cache object describes the cache control parameters to be applied to the content by intermediate caches.

Property: ignore-query-string

Description: Allows a cache to ignore URI query string parameters while comparing URIs for equivalence. Each query parameter to ignore is specified in the list. If all query parameters should be ignored, then the list MUST be empty.

Type: List of String

Mandatory-to-Specify: No. Default is to consider query string parameters when comparing URIs or to rely on other properties of the Cache object.

4.2.8. Grouping

A Grouping object identifies a large group of content to which this content belongs.

Property: ccid

Description: Content Collection identifier for an application-specific purpose such as logging.

Type: String

Mandatory-to-Specify: No. Default is an empty string.

Property: sid

Description: Session identifier for an application-specific purpose such as logging.

Type: String

Mandatory-to-Specify: No. Default is an empty string.

4.3. CDNI Metadata Simple Data Type Descriptions

This section describes the simpler data types that are used for properties of CDNI Metadata objects.

4.3.1. Link

A link object may be used in place of any of the objects or properties described above. Links can be used to avoid duplication if the same metadata information is repeated within the metadata tree. When a link replaces an object, its href property is set to the URI of the resource, its rel property is set to the name of the property it is replacing, and its type property is set to the type of the object it is replacing.

Property: href

Description: The URI of the of the addressable object being referenced.

Type: URI

Mandatory-to-Specify: Yes

Property: rel

Description: The Relationship between the referring object and the object it is referencing.

Type: String

Mandatory-to-Specify: Yes

Property: type

Description: The type of the object being referenced.

Type: String

Mandatory-to-Specify: Yes

4.3.2. Protocol

Protocol objects are used to specify registered protocols for content acquisition or delivery (see Section 7.1.1.2).

Type: String

Mandatory-to-Specify: Yes

4.3.3. Endpoint

A hostname (with optional port) or an IP address (with optional port).

Note: All implementations MUST support IPv4 addresses encoded as specified by the 'IPv4address' rule in Section 3.2.2 of [RFC3986] and MUST support all IPv6 address formats specified in [RFC4291]. Server implementations SHOULD use IPv6 address formats specified in [RFC5952].

4.3.4. URI

A URI as specified in [RFC3986].

4.3.5. Time

A time value expressed in seconds since Unix epoch in the UTC timezone.

5. CDNI Metadata Capabilities

CDNI Metadata is used to convey information pertaining to content delivery from uCDN to dCDN. For optional metadata, it may be useful for the uCDN to know if the dCDN supports the metadata, prior to delegating any content requests to the dCDN. If optional-to-implement metadata is mandatory-to-enforce and the dCDN does not support it, any delegated requests for that content will fail, so there is no reason to delegate those requests. Likewise, for any metadata which may be assigned optional values, it may be useful for the uCDN to know which values the dCDN supports, prior to delegating any content requests to the dCDN. If a the optional value assigned to a given piece of content's metadata is not supported by the dCDN, any delegated requests for that content may fail, so there is likely no reason to delegate those requests.

The CDNI Footprint and Capabilities Interface provides a means of advertising capabilities from dCDN to uCDN. Support for optional metadata and support for optional metadata values may be advertised using the capabilities interface. This section describes the capabilities advertisement requirements for the metadata defined in Section 4.2

5.1. Protocol ACL Capabilities

The ProtocolACL object contains a list of Protocol values. The dCDN MUST advertise which delivery protocols it supports so that the uCDN knows what type of content requests it can redirect to the dCDN. If the dCDN does not support a given acquisition or delivery protocol, the uCDN should not delegate requests requiring those protocols to the dCDN as the dCDN will not be able to properly acquire or deliver the content.

ProtocolRules are defined for either acquisition or delivery. For some CDNs, certain combinations of acquisition and delivery protocols may not make sense (e.g., RTSP acquisition for HTTP delivery), while other CDNs may support customized protocol adaptation. ProtocolACL capabilities are not intended to define which combinations of protocols should be used. ProtocolACL capabilities are only intended to describe which protocols the dCDN does or does not support. Protocol combination restrictions are specified in the metadata itself and associated with specific groups of content assets.

5.2. Authorization Metadata Capabilities

The Authorization object contains a list of Auth values. The dCDN MUST advertise which authorization algorithms it supports so that the uCDN knows what type of content requests it can redirect to the dCDN. If the dCDN does not support a given authorization algorithm, the uCDN should not delegate requests requiring that algorithm to the dCDN as the dCDN will not be able to properly acquire the content or enforce delivery restrictions.

6. CDNI Metadata interface

This section specifies an interface to enable a Downstream CDN to retrieve CDNI Metadata objects from an Upstream CDN.

The interface can be used by a Downstream CDN to retrieve CDNI Metadata objects either dynamically as required by the Downstream CDN to process received requests (for example in response to receiving a CDNI Request Routing request from an Upstream CDN or in response to receiving a request for content from a User Agent) or in advance of being required (for example in case of prepositioned CDNI Metadata acquisition).

The CDNI Metadata interface is built on the principles of RESTful web services. This means that requests and responses over the interface are built around the transfer of representations of hyperlinked resources. A resource in the context of the CDNI Metadata interface is any object in the Data Model (as described in Section 3 through Section 4).

In the general case a CDNI Metadata server makes each instance of an addressable CDNI Metadata object available via a unique URI that returns a representation of that instance of that CDNI Metadata object. When an object needs to reference another addressable CDNI Metadata object (for example a HostIndex object referencing a HostMetadata object) it does so by including a link to the referenced object.

CDNI Metadata servers are free to assign whatever structure they desire to the URIs for CDNI Metadata objects and CDNI Metadata clients MUST NOT make any assumptions regarding the structure of CDNI Metadata URIs or the mapping between CDNI Metadata objects and their associated URIs. Therefore any URIs present in the examples below are purely illustrative and are not intended to impose a definitive structure on CDNI Metadata interface implementations.

6.1. Transport

The CDNI Metadata interface uses HTTP as the underlying protocol transport.

The HTTP Method in the request defines the operation the request would like to perform. Servers implementing the CDNI Metadata interface MUST support the HTTP GET and HEAD methods.

The corresponding HTTP Response returns the status of the operation in the HTTP Status Code and returns the current representation of the resource (if appropriate) in the Response Body. HTTP Responses from servers implementing the CDNI Metadata interface that contain a response body SHOULD include an ETag to enable validation of cached versions of returned resources.

The CDNI Metadata interface specified in this document is a read-only interface. Therefore support for other HTTP methods such as PUT, POST and DELETE etc. is not specified. Server implementations of this interface SHOULD reject all methods other than GET and HEAD.

As the CDNI Metadata interface builds on top of HTTP, CDNI Metadata servers may make use of any HTTP feature when implementing the CDNI Metadata interface, for example a CDNI Metadata server may make use of HTTP's caching mechanisms to indicate that the returned response/representation can be reused without re-contacting the CDNI Metadata server.

6.2. Retrieval of CDNI Metadata resources

In the general case a CDNI Metadata server makes each instance of an addressable CDNI Metadata object available via a unique URI and therefore in order to retrieve CDNI Metadata, a CDNI Metadata client first makes a HTTP GET request for the URI of the HostIndex which provides the CDNI Metadata client with a list of Hostnames for which the upstream CDN may delegate content delivery to the downstream CDN.

In order to retrieve the CDNI Metadata for a particular request the CDNI Metadata client processes the received HostIndex object and finds the corresponding HostMetadata entry (by matching the hostname in the request against the hostnames in the HostMatch). If the HostMetadata is linked (rather than embedded), the CDNI metadata client then makes a GET request for the URI specified in the href property of the Link object which points to the HostMetadata object itself.

In order to retrieve the most specific metadata for a particular request, the CDNI metadata client inspects the HostMetadata for references to more specific PathMetadata objects. If any PathMetadata match the request (and are linked rather than embedded), the CDNI metadata client makes another GET request for the PathMetadata. Each PathMetadata object may also include references to yet more specific metadata. If this is the case, the CDNI metadata client continues requesting PathMetadata recursively.

Where a downstream CDN is interconnected with multiple upstream CDNs, the downstream CDN must decide which upstream CDN's CDNI metadata should be used to handle a particular User Agent request.

When application level redirection (e.g. HTTP 302 redirects) is being used between CDNs, it is expected that the downstream CDN will be able to determine the upstream CDN that redirected a particular request from information contained in the received request (e.g. via the URI). With knowledge of which upstream CDN routed the request, the downstream CDN can choose the correct metadata server from which to obtain the HostIndex. Note that the HostIndex served by each uCDN may be unique.

In the case of DNS redirection there is not always sufficient information carried in the DNS request from User Agents to determine the upstream CDN that redirected a particular request (e.g. when content from a given host is redirected to a given downstream CDN by more than one upstream CDN) and therefore downstream CDNs may have to apply local policy when deciding which upstream CDN's metadata to apply.

6.3. Bootstrapping

The URI for the HostIndex object of a given upstream CDN needs to be either discovered by or configured in the downstream CDN. All other objects/resources are then discoverable from the HostIndex object by following the links in the HostIndex object and the referenced HostMetadata and PathMetadata objects.

If the URI for the HostIndex object is not manually configured in the downstream CDN then the HostIndex URI could be discovered. A mechanism allowing the downstream CDN to discover the URI of the HostIndex is outside the scope of this document.

6.4. Encoding

Object are resources that may be:

- o Addressable, where the object is a resource that may be retrieved or referenced via its own URI.
- o Embedded, where the object is contained (or inlined) within a property of an addressable object.

In the descriptions of objects we use the term "X contains Y" to mean either Y is directly embedded in X or that Y is linked to by X. It is generally a deployment choice for the uCDN implementation to decide when and which CDNI Metadata objects to embed and which are separately addressable.

6.4.1. MIME Media Types

All MIME types are prefixed with "application/cdni." The MIME type for each object matches the type name of that object as defined by this document. Table 3 lists a few examples of the MIME Media Type for each object (resource) that is retrievable through the CDNI Metadata interface.

Data Object	MIME Media Type
HostIndex	application/cdni.HostIndex
HostMatch	application/cdni.HostMatch
HostMetadata	application/cdni.HostMetadata
PathMatch	application/cdni.PathMatch
PathMetadata	application/cdni.PathMetadata

Table 3: Example MIME Media Types for CDNI Metadata objects

See <http://www.iana.org/assignments/media-types/index.html> for reference.

6.4.2. JSON Encoding of Objects

CDNI Metadata objects are encoded as JSON objects containing a dictionary of (key,value) pairs where the keys are the property names and the values are the associated property values.

The keys of the dictionary are the names of the properties associated with the object and are therefore dependent on the specific object being encoded (i.e. dependent on the MIME Media Type of the returned

resource). Likewise, the values associated with each key are dependent on the specific object being encoded (i.e. dependent on the MIME Media Type of the returned resource).

Dictionary keys in JSON are case sensitive and therefore by convention any dictionary key defined by this document (for example the names of CDNI Metadata object properties) MUST be represented in lowercase.

In addition to the properties specific to each object type, the keys defined below may be present in any object.

Key: base

Description: Provides a prefix for any relative URLs in the object. This is similar to the XML base tag [XML-BASE]. If absent, all URLs in the remainder of the document must be absolute URLs.

Type: URI

Mandatory: No

Key: links

Description: The links of this object to other addressable objects. Any property may be replaced by a link to an object with the same type as the property it replaces.

Type: List of Link objects

Mandatory: Yes

6.4.2.1. JSON Example

A downstream CDN may request the HostIndex and receive the following object of type "application/cdni.HostIndex+json":

```
{
  "hosts": [
    {
      "host": "video.example.com",
      "links": [
        {
          "rel": "host-metadata",
          "type": "application/cdni.HostMetadata",
          "href": "http://metadata.example.ucdn.com/video"
        }
      ]
    }
  ]
}
```

```

    ]
  },
  {
    "host": "images.example.com",
    "links": [
      {
        "rel": "host-metadata",
        "type": "application/cdni.HostMetadata",
        "href": "http://metadata.ucdn.example.com/images"
      }
    ]
  }
]
}

```

If the incoming request has a Host header with "video.example.com" then the downstream CDN would fetch from the next metadata object from "http://metadata.ucdn.example.com/video" expecting a MIME type of "application/cdni.HostMetadata+json":

```

{
  "metadata": [
    {
      "type": "application/cdni.SourceMetadata",
      "value": {
        "sources": [
          {
            "links": [{
              "rel": "auth",
              "type": "application/cdni.Auth",
              "href": "http://metadata.ucdn.example.com/auth1234"
            }],
            "endpoint": "acq1.ucdn.example.com",
            "protocol": "ftp"
          },
          {
            "links": [{
              "rel": "auth",
              "type": "application/cdni.Auth",
              "href": "http://metadata.ucdn.example.com/auth1234"
            }],
            "endpoint": "acq2.ucdn.example.com",
            "protocol": "http"
          }
        ]
      }
    }
  ],
}

```

```
{
  "type": "application/cdni.LocationACL",
  "value": {
    "locations": [
      {
        "locations": [
          { "iprange": "192.168.0.0/16" }
        ],
        "action": "deny"
      }
    ]
  }
},
{
  "type": "application/cdni.ProtocolACL",
  "value": {
    "protocols": [
      {
        "protocols": [
          "ftp"
        ],
        "action": "deny"
      }
    ]
  }
},
],
"paths": [
  {
    "path-pattern": {
      "pattern": "/videos/trailers/*"
    },
    "links": [{
      "rel": "path-metadata",
      "type": "application/cdni.PathMetadata",
      "href": "http://metadata.ucdn.example.com/videos/trailers"
    }]
  },
  {
    "path-pattern": {
      "pattern": "/videos/movies/*"
    },
    "links": [{
      "rel": "pathmetadata",
      "type": "application/cdni.PathMetadata",
      "href": "http://metadata.ucdn.example.com/videos/movies"
    }]
  }
]
```

```
]
}
```

Suppose the path of the requested resource matches the `"/video/movies/*"` pattern, the next metadata requested would be for `"http://metadata.ucdn.example.com/video/movies"` with an expected type of `"application/cdni.PathMetadata"`:

```
{
  "metadata": [],
  "paths": [
    {
      "path-pattern": {
        "pattern": "/videos/movies/hd/*"
      },
      "links": [{
        "rel": "pathmetadata",
        "type": "application/cdni.PathMetadata",
        "href": "http://metadata.ucdn.example.com/videos/movies/hd"
      }]
    }
  ]
}
```

Finally, if the path of the requested resource also matches the `"/videos/movies/hd/*"` pattern, the downstream CDN would also fetch the following object from `"http://metadata.ucdn.example.com/videos/movies/hd"` with MIME type `"application/cdni.PathMetadata"`:

```
{
  "metadata": [
    {
      "type": "application/cdni.TimeWindowACL",
      "value": {
        "times": [
          {
            "start": "1213948800",
            "end": "1327393200"
          }
        ],
        "type": "allow"
      }
    }
  ]
}
```

```
}
```

6.5. Extensibility

The set of property Metadata may be extended with proprietary and/or custom property Metadata. The GenericMetadata object defined in Section 4.1.7 allows any Metadata property to be included in either the HostMetadata or PathMetadata lists. As described in Section 3.4, any proprietary and/or custom property Metadata SHOULD be identified by the "ext." prefix in an appropriately descriptive type which conveys the organization defining the property Metadata and the function of the property Metadata.

Note: Identification of the property Metadata defining organization in the property Metadata type decreases the possibility of property Metadata type collision. The fully-qualified domain name of the organization in reverse order may be used for this purpose.

6.5.1. Metadata Enforcement

At any given time, the set of property Metadata supported by the uCDN may not match the set of property Metadata supported by the dCDN. The uCDN may or may not know which property Metadata the dCDN supports. In cases where the uCDN supports Metadata that the dCDN does not, the dCDN MUST be aware of any Metadata marked as "mandatory-to-enforce". If a CDN does not understand or is unable to perform the functions associated with any "mandatory-to-enforce" Metadata, the CDN MUST NOT service any requests for the corresponding content.

Any standard which defines a new GenericMetadata Type SHOULD also define whether or not the new metadata is mandatory-to-enforce.

Note: Ideally, uCDNs would not delegate content requests to a dCDN which does not support the mandatory-to-enforce Metadata associated with the content being requested. However, even if the uCDN has a priori knowledge of the Metadata supported by the dCDN (e.g., via the CDNI capabilities interface or through out-of-band negotiation between CDN operators) Metadata support may fluctuate or be inconsistent (e.g., due to mis-communication, mis-configuration, or temporary outage). Thus, the dCDN MUST evaluate all Metadata associated with content requests and reject any requests where "mandatory-to-enforce" Metadata associated with the content cannot be enforced.

6.5.2. Metadata Override

It is possible that new Metadata definitions may obsolete or override existing property Metadata (e.g., a future revision of the CDNI Metadata interface may redefine the Auth Metadata or a custom vendor extension may implement an alternate Auth Metadata option). If multiple Metadata (e.g., cdni.v2.Auth, ext.vendor1.Auth, and ext.vendor2.Auth) all override an existing Metadata (e.g., cdni.Auth) and all are marked as "mandatory-to-enforce", it may be ambiguous which Metadata should be applied, especially if the functionality of the Metadata conflict.

As described in Section 3.3, Metadata override only applies to Metadata objects of the same exact type, found in HostMetadata and nested PathMetadata structures. The CDNI Metadata interface does not support enforcement of dependencies between different Metadata types. It is the responsibility of the CSP and the CDN operators to ensure that Metadata assigned to a given content do not conflict.

Note: Because Metadata is inherently ordered in GenericMetadata lists, as well as in the PathMetadata hierarchy and PathMatch lists, multiple conflicting Metadata types MAY be used, however, Metadata hierarchies MUST ensure that independent PathMatch root objects are used to prevent ambiguous or conflicting Metadata definitions.

6.6. Versioning

The version of CDNI Metadata Structural objects is specified by the HTTP Content-Type header. Upon responding to a request for an object, a metadata server MUST include a Content-Type header with the MIME-type and version number of the object. HTTP requests sent to a metadata server SHOULD include an Accept header with the MIME-type and version of the expected object. Unless stated otherwise, the version of each object defined by this document is version 1. For example: "Content-Type: application/cdni.HostIndex.v1":.

GenericMetadata objects include a "type" property which specifies the MIME-type of the GenericMetadata value. This MIME-type should also include a version. Any document which defines a new type of GenericMetadata should specify the version number which it describes. For example: "application/cdni.Location.v1".

7. IANA Considerations

This document requests the registration of the "application/cdni" MIME Media Type under the IANA MIME Media Type registry (<http://www.iana.org/assignments/media-types/index.html>).

7.1. GenericMetadata Type Registry

CDNI Metadata is distributed as a list of GenericMetadata objects which specify a type field and a type-specific value field, as described in Section 4.1.7. In order to prevent namespace collisions for GenericMetadata object types a new IANA registry is requested for "CDNI GenericMetadata Types" namespace. The namespace shall be split into two partitions: standard and vendor defined. As described in Section 4.1.7 and Section 6.5, the vendor defined namespace partition SHOULD use a namespace prefix of "ext.", while the standard namespace partition MUST NOT.

The standard namespace partition MUST conform to the "RFC Required" policy as defined in [RFC5226]. The vendor defined namespace partition should be further partitioned into vendor specific partitions with the prefix "ext.vendor_name.". The vendor defined partition SHOULD conform to the "Expert Review" policy as defined in [RFC5226]. The expert review is simply to prevent namespace hoarding. The vendor specific partitions MAY conform to the "First Come First Served" policy as defined in [RFC5226], however, vendors defining new GenericMetadata types which conflict with other GenericMetadata types SHOULD follow the guidelines for the "Specification Required" policy as defined in [RFC5226].

The following table defines the initial values for the standard partition:

type name	Specification
SourceMetadata	RFCThis
LocationACL	RFCThis
TimeWindowACL	RFCThis
ProtocolACL	RFCThis
Auth	RFCThis
Cache	RFCThis
Grouping	RFCThis

7.1.1. GenericMetadata Sub-Registries

Some of the initial standard GenericMetadata objects contain enumerated types which require registration (i.e., LocationACL footprint types, ProtocolACL protocols, and Auth protocols). The following sections define the initial values for these GenericMetadata type sub-registries.

7.1.1.1. Footprint Sub-Registry

The "CDNI Metadata Footprint Types" namespace defines the valid Footprint object type values used by the Footprint object in Section 4.2.2.2. Additions to the Footprint type namespace MUST conform to the "Expert Review" policy as defined in [RFC5226]. The expert review should verify that new type definitions do not duplicate existing type definitions and prevent gratuitous additions to the namespace.

The following table defines the initial Footprint type values:

type name	definition
IPv4	Single IPv4 address
IPv6	Single IPv6 address
IPv4Range	List of contiguous IPv4 addresses denoted by a start address and an end address separated by a dash (e.g., 192.168.0.1-192.168.0.20), inclusive.
IPv6Range	List of contiguous IPv6 addresses denoted by a start address and an end address separated by a dash (e.g., fc80::0001-fc80::0014), inclusive.
IPv4CIDR	IPv4 address block using slash prefix length notation (e.g., 192.168.0.16/28).
IPv6CIDR	IPv6 address block using slash prefix length notation (e.g., fc80::0010/124).
ASN	Autonomous System (AS) Number
CountryCode	ISO 3166-1 alpha-2 code
DVDRegion	DVD Region code (i.e., integer in the range 0-6).

7.1.1.2. Protocol Sub-Registry

The "CDNI Metadata Protocols" namespace defines the valid Protocol object values in Section 4.3.2, used by the SourceMetadata and ProtocolACL objects. Additions to the Protocol namespace MUST conform to the "Expert Review" policy as defined in [RFC5226]. The expert review should verify that new type definitions do not duplicate existing type definitions and prevent gratuitous additions to the namespace.

The following table defines the initial Protocol values:

protocol name	definition
HTTP	TBD
HTTPS	TBD

RTSP	TBD	
RTMP	TBD	
+-----+	+-----+	+-----+

7.1.1.3. Authentication Sub-Registry

The "CDNI Metadata Auth" namespace defines the valid Auth object types used by the Auth object in Section 4.2.6. Additions to the Footprint type namespace MUST conform to the "Expert Review" policy as defined in [RFC5226]. The expert review should verify that new type definitions do not duplicate existing type definitions and prevent gratuitous additions to the namespace.

The following table defines the initial Auth type values:

type name	definition
credentials	Simple username and password authentication as defined by Section 4.2.6.1.

8. Security Considerations

The CDNI Metadata Interface is expected to be secured as a function of the transport protocol (e.g. HTTP authentication [RFC2617], HTTPS [RFC2818], or inter-domain IPSec).

If a malicious metadata server is contacted by a downstream CDN, the malicious server may provide metadata to the downstream CDN which denies service for any piece of content to any user agent. The malicious server may also provide metadata which directs a downstream CDN to a malicious origin server instead of the actual origin server. The dCDN is expected to authenticate the server to prevent this situation (e.g. by using HTTPS and validating the server's certificate).

A malicious metadata client could request metadata for a piece of content from an upstream CDN. The metadata information may then be used to glean information regarding the uCDN or to contact an upstream origin server. The uCDN is expected to authenticate client requests to prevent this situation.

9. Acknowledgements

The authors would like to thank David Ferguson and Francois le Faucheur for their valuable comments and input to this document.

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2617] Franks, J., Hallam-Baker, P., Hostetler, J., Lawrence, S., Leach, P., Luotonen, A., and L. Stewart, "HTTP Authentication: Basic and Digest Access Authentication", RFC 2617, June 1999.
- [RFC2818] Rescorla, E., "HTTP Over TLS", RFC 2818, May 2000.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [RFC5952] Kawamura, S. and M. Kawashima, "A Recommendation for IPv6 Address Text Representation", RFC 5952, August 2010.

10.2. Informative References

- [I-D.ietf-cdni-framework]
Peterson, L. and B. Davie, "Framework for CDN Interconnection", draft-ietf-cdni-framework-03 (work in progress), February 2013.
- [I-D.ietf-cdni-requirements]
Leung, K. and Y. Lee, "Content Distribution Network Interconnection (CDNI) Requirements", draft-ietf-cdni-requirements-05 (work in progress), February 2013.
- [I-D.zyp-json-schema]
Zyp, K. and G. Court, "A JSON Media Type for Describing the Structure and Meaning of JSON Documents", draft-zyp-json-schema-03 (work in progress), November 2010.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005.

- [RFC4151] Kindberg, T. and S. Hawke, "The 'tag' URI Scheme", RFC 4151, October 2005.
- [RFC4287] Nottingham, M., Ed. and R. Sayre, Ed., "The Atom Syndication Format", RFC 4287, December 2005.
- [RFC6707] Niven-Jenkins, B., Le Faucheur, F., and N. Bitar, "Content Distribution Network Interconnection (CDNI) Problem Statement", RFC 6707, September 2012.
- [XML-BASE] Marsh, J., Ed. and R. Tobin, Ed., "XML Base (Second Edition) - <http://www.w3.org/TR/xmlbase/>", January 2009.

Appendix A. Relationship to the CDNI Requirements

Section 6 of [I-D.ietf-cdni-requirements] lists the requirements for the CDNI Metadata Distribution interface. This section outlines which of those requirements are met by the CDNI Metadata interface specified in this document.

All metadata requirements are met either directly or indirectly by the CDNI Metadata Interface described in this document, with the clarifications or exceptions described in the following paragraphs.

Requirements related to pre-positioning of metadata are met by this document on the assumption that other CDNI Interfaces are to be used by the upstream CDN to trigger the pre-positioning of metadata by the downstream CDN via the CDNI Metadata Interface. Triggering metadata pre-positioning is beyond the scope of the CDNI Metadata interface. However, the interface as described by this document supports pulling metadata on-demand for the purpose of pre-positioning.

Requirement META-7 relating to modification of metadata by the upstream CDN is met both by allowing timeouts on the cacheability of metadata objects and by allowing other CDNI interfaces to initiate a refetch or purge of metadata.

Requirement META-18 relating to surrogate cache behavior parameters is supported via extensibility. However, the example parameters in META-18 are not described in this document.

Authors' Addresses

Ben Niven-Jenkins
Velocix (Alcatel-Lucent)
3 Ely Road
Milton, Cambridge CB24 6AA
UK

Email: ben@velocix.com

Rob Murray
Velocix (Alcatel-Lucent)
3 Ely Road
Milton, Cambridge CB24 6AA
UK

Email: rmurray@velocix.com

Grant Watson
Velocix (Alcatel-Lucent)
3 Ely Road
Milton, Cambridge CB24 6AA
UK

Email: gwatson@velocix.com

Matt Caulfield
Cisco Systems
1414 Massachusetts Avenue
Boxborough, MA 01719
USA

Phone: +1 978 936 9307
Email: mcaulfie@cisco.com

Kent Leung
Cisco Systems
3625 Cisco Way
San Jose 95134
USA

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

Kevin J. Ma
Azuki Systems, Inc.
43 Nagog Park
Acton, MA 01720
USA

Phone: +1 978-844-5100
Email: kevin.ma@azukisystems.com

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: March 1, 2017

B. Niven-Jenkins
R. Murray
Velocix (Alcatel-Lucent)
M. Caulfield
Cisco Systems
K. Ma
Ericsson
August 28, 2016

CDN Interconnection Metadata
draft-ietf-cdni-metadata-21

Abstract

The Content Delivery Networks Interconnection (CDNI) metadata interface enables interconnected Content Delivery Networks (CDNs) to exchange content distribution metadata in order to enable content acquisition and delivery. The CDNI metadata associated with a piece of content provides a downstream CDN with sufficient information for the downstream CDN to service content requests on behalf of an upstream CDN. This document describes both a base set of CDNI metadata and the protocol for exchanging that metadata.

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 March 1, 2017.

Copyright Notice

Copyright (c) 2016 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	5
1.2. Supported Metadata Capabilities	5
2. Design Principles	6
3. CDNI Metadata object model	7
3.1. HostIndex, HostMatch, HostMetadata, PathMatch, PatternMatch and PathMetadata objects	8
3.2. Generic CDNI Metadata Objects	10
3.3. Metadata Inheritance and Override	13
4. CDNI Metadata objects	14
4.1. Definitions of the CDNI structural metadata objects	15
4.1.1. HostIndex	15
4.1.2. HostMatch	15
4.1.3. HostMetadata	17
4.1.4. PathMatch	18
4.1.5. PatternMatch	19
4.1.6. PathMetadata	20
4.1.7. GenericMetadata	21
4.2. Definitions of the initial set of CDNI Generic Metadata objects	23
4.2.1. SourceMetadata	23
4.2.1.1. Source	24
4.2.2. LocationACL Metadata	25
4.2.2.1. LocationRule	27
4.2.2.2. Footprint	27
4.2.3. TimeWindowACL	29
4.2.3.1. TimeWindowRule	30
4.2.3.2. TimeWindow	31
4.2.4. ProtocolACL Metadata	31
4.2.4.1. ProtocolRule	32
4.2.5. DeliveryAuthorization Metadata	33

4.2.6.	Cache	34
4.2.7.	Auth	36
4.2.8.	Grouping	37
4.3.	CDNI Metadata Simple Data Type Descriptions	37
4.3.1.	Link	37
4.3.1.1.	Link Loop Prevention	39
4.3.2.	Protocol	39
4.3.3.	Endpoint	39
4.3.4.	Time	40
4.3.5.	IPv4CIDR	40
4.3.6.	IPv6CIDR	40
4.3.7.	ASN	41
4.3.8.	CountryCode	41
5.	CDNI Metadata Capabilities	41
6.	CDNI Metadata interface	42
6.1.	Transport	42
6.2.	Retrieval of CDNI Metadata resources	43
6.3.	Bootstrapping	44
6.4.	Encoding	44
6.5.	Extensibility	45
6.6.	Metadata Enforcement	46
6.7.	Metadata Conflicts	46
6.8.	Versioning	47
6.9.	Media Types	48
6.10.	Complete CDNI Metadata Example	48
7.	IANA Considerations	52
7.1.	CDNI Payload Types	52
7.1.1.	CDNI MI HostIndex Payload Type	53
7.1.2.	CDNI MI HostMatch Payload Type	53
7.1.3.	CDNI MI HostMetadata Payload Type	54
7.1.4.	CDNI MI PathMatch Payload Type	54
7.1.5.	CDNI MI PatternMatch Payload Type	54
7.1.6.	CDNI MI PathMetadata Payload Type	54
7.1.7.	CDNI MI SourceMetadata Payload Type	54
7.1.8.	CDNI MI Source Payload Type	55
7.1.9.	CDNI MI LocationACL Payload Type	55
7.1.10.	CDNI MI LocationRule Payload Type	55
7.1.11.	CDNI MI Footprint Payload Type	55
7.1.12.	CDNI MI TimeWindowACL Payload Type	55
7.1.13.	CDNI MI TimeWindowRule Payload Type	56
7.1.14.	CDNI MI TimeWindow Payload Type	56
7.1.15.	CDNI MI ProtocolACL Payload Type	56
7.1.16.	CDNI MI ProtocolRule Payload Type	56
7.1.17.	CDNI MI DeliveryAuthorization Payload Type	56
7.1.18.	CDNI MI Cache Payload Type	57
7.1.19.	CDNI MI Auth Payload Type	57
7.1.20.	CDNI MI Grouping Payload Type	57
7.2.	CDNI Metadata Footprint Types Registry	57

7.3. CDNI Metadata Protocol Types Registry	58
8. Security Considerations	58
8.1. Authentication and Integrity	59
8.2. Confidentiality and Privacy	59
8.3. Securing the CDNI Metadata interface	60
9. Acknowledgements	60
10. Contributing Authors	60
11. References	61
11.1. Normative References	61
11.2. Informative References	63
Authors' Addresses	64

1. Introduction

Content Delivery Networks Interconnection (CDNI) [RFC6707] enables a downstream Content Delivery Network (dCDN) to service content requests on behalf of an upstream CDN (uCDN).

The CDNI metadata interface is discussed in [RFC7336] along with four other interfaces that can be used to compose a CDNI solution (CDNI Control interface, CDNI Request Routing Redirection interface, CDNI Footprint & Capabilities Advertisement interface and CDNI Logging interface). [RFC7336] describes each interface and the relationships between them. The requirements for the CDNI metadata interface are specified in [RFC7337].

The CDNI metadata associated with a piece of content (or with a set of content) provides a dCDN with sufficient information for servicing content requests on behalf of an uCDN, in accordance with the policies defined by the uCDN.

This document defines the CDNI metadata interface which enables a dCDN to obtain CDNI metadata from an uCDN so that the dCDN can properly process and respond to:

- o Redirection requests received over the CDNI Request Routing Redirection interface [I-D.ietf-cdni-redirection].
- o Content requests received directly from User Agents.

Specifically, this document specifies:

- o A data structure for mapping content requests and redirection requests to CDNI metadata objects (Section 3 and Section 4.1).
- o An initial set of CDNI Generic metadata objects (Section 4.2).
- o A HTTP web service for the transfer of CDNI metadata (Section 6).

1.1. Terminology

This document reuses the terminology defined in [RFC6707].

Additionally, the following terms are used throughout this document and are defined as follows:

- o Object - a collection of properties.
- o Property - a key and value pair where the key is a property name and the value is the property value or another object.

This document uses the phrase "[Object] A contains [Object] B" for simplicity when a strictly accurate phrase would be "[Object] A contains or references (via a Link object) [Object] B".

1.2. Supported Metadata Capabilities

Only the metadata for a small set of initial capabilities is specified in this document. This set provides the minimum amount of metadata for basic CDN interoperability while still meeting the requirements set forth by [RFC7337].

The following high-level functionality can be configured via the CDNI metadata objects specified in Section 4:

- o Acquisition Source: Metadata for allowing a dCDN to fetch content from a uCDN.
- o Delivery Access Control: Metadata for restricting (or permitting) access to content based on any of the following factors:
 - * Location
 - * Time Window
 - * Delivery Protocol
- o Delivery Authorization: Metadata for authorizing dCDN user agent requests.
- o Cache Control: Metadata for controlling cache behavior of the dCDN.

The metadata encoding described by this document is extensible in order to allow for future additions to this list.

The set of metadata specified in this document covers the initial capabilities above. It is only intended to support CDN interconnection for the delivery of content by a dCDN using HTTP/1.1 [RFC7230] and for a dCDN to be able to acquire content from a uCDN using either HTTP/1.1 or HTTP/1.1 over TLS [RFC2818].

Supporting CDN interconnection for the delivery of content using unencrypted HTTP/2 [RFC7540] (as well as for a dCDN to acquire content using unencrypted HTTP/2 or HTTP/2 over TLS) requires the registration of these protocol names in the CDNI Metadata Protocol Types registry Section 7.3.

Delivery of content using HTTP/1.1 over TLS or HTTP/2 over TLS SHOULD follow the guidelines set forth in [RFC7525]. Offline configuration of TLS parameters between CDNs is beyond the scope of this document.

2. Design Principles

The CDNI metadata interface was designed to achieve the following objectives:

1. Cacheability of CDNI metadata objects;
2. Deterministic mapping from redirection requests and content requests to CDNI metadata properties;
3. Support for DNS redirection as well as application-specific redirection (for example HTTP redirection);
4. Minimal duplication of CDNI metadata; and
5. Leveraging of existing protocols.

Cacheability can decrease the latency of acquiring metadata while maintaining its freshness, and therefore decrease the latency of serving content requests and redirection requests, without sacrificing accuracy. The CDNI metadata interface uses HTTP and its existing caching mechanisms to achieve CDNI metadata cacheability.

Deterministic mappings from content to metadata properties eliminates ambiguity and ensures that policies are applied consistently by all dCDNs.

Support for both HTTP and DNS redirection ensures that the CDNI metadata meets the same design principles for both HTTP and DNS based redirection schemes.

Minimal duplication of CDNI metadata improves storage efficiency in the CDNs.

Leveraging existing protocols avoids reinventing common mechanisms such as data structure encoding (by leveraging I-JSON [RFC7493]) and data transport (by leveraging HTTP [RFC7230]).

3. CDNI Metadata object model

The CDNI metadata object model describes a data structure for mapping redirection requests and content requests to metadata properties. Metadata properties describe how to acquire content from a uCDN, authorize access to content, and deliver content from a dCDN. The object model relies on the assumption that these metadata properties can be grouped based on the hostname of the content and subsequently on the resource path (URI) of the content. The object model associates a set of CDNI metadata properties with a Hostname to form a default set of metadata properties for content delivered on behalf of that Hostname. That default set of metadata properties can be overridden by properties that apply to specific paths within a URI.

Different Hostnames and URI paths will be associated with different sets of CDNI metadata properties in order to describe the required behaviour when a dCDN surrogate or request router is processing User Agent requests for content at that Hostname and URI path. As a result of this structure, significant commonality could exist between the CDNI metadata properties specified for different Hostnames, different URI paths within a Hostname and different URI paths on different Hostnames. For example the definition of which User Agent IP addresses should be grouped together into a single network or geographic location is likely to be common for a number of different Hostnames; although a uCDN is likely to have several different policies configured to express geo-blocking rules, it is likely that a single geo-blocking policy could be applied to multiple Hostnames delivered through the CDN.

In order to enable the CDNI metadata for a given Hostname and URI Path to be decomposed into reusable sets of CDNI metadata properties, the CDNI metadata interface splits the CDNI metadata into separate objects. Efficiency is improved by enabling a single CDNI metadata object (that is shared across Hostname and/or URI paths) to be retrieved and stored by a dCDN once, even if it is referenced by the CDNI metadata for multiple Hostnames and/or URI paths.

Important Note: Any CDNI metadata object A that contains another CDNI metadata object B can include a Link object specifying a URI that can be used to retrieve object B, instead of embedding object B within object A. The remainder of this document uses the phrase "[Object] A

contains [Object] B" for simplicity when a strictly accurate phrase would be "[Object] A contains or references (via a Link object) [Object] B". It is generally a deployment choice for the uCDN implementation to decide when to embed CDNI metadata objects and when to reference separate resources via Link objects.

Section 3.1 introduces a high level description of the HostIndex, HostMatch, HostMetadata, PathMatch, PatternMatch and PathMetadata objects, and describes the relationships between them.

Section 3.2 introduces a high level description of the CDNI GenericMetadata object which represents the level at which CDNI metadata override occurs between HostMetadata and PathMetadata objects.

Section 4 describes in detail the specific CDNI metadata objects and properties specified by this document which can be contained within a CDNI GenericMetadata object.

3.1. HostIndex, HostMatch, HostMetadata, PathMatch, PatternMatch and PathMetadata objects

The relationships between the HostIndex, HostMatch, HostMetadata, PathMatch, PatternMatch and PathMetadata objects are described in Figure 1.




Figure 1: Relationships between CDNI Metadata Objects (Diagram Representation)

A HostIndex object (see Section 4.1.1) contains an array of HostMatch objects (see Section 4.1.2) that contain Hostnames (and/or IP addresses) for which content requests might be delegated to the dCDN. The HostIndex is the starting point for accessing the uCDN CDNI metadata data store. It enables the dCDN to deterministically discover which CDNI metadata objects it requires in order to deliver a given piece of content.

The HostIndex links Hostnames (and/or IP addresses) to HostMetadata objects (see Section 4.1.3) via HostMatch objects. A HostMatch object defines a Hostname (or IP address) to match against a requested host and contains a HostMetadata object.

HostMetadata objects contain the default GenericMetadata objects (see Section 4.1.7) required to serve content for that host. When looking up CDNI metadata, the dCDN looks up the requested Hostname (or IP address) against the HostMatch entries in the HostIndex, from there it can find HostMetadata which describes the default metadata properties for each host as well as PathMetadata objects (see Section 4.1.6), via PathMatch objects (see Section 4.1.4). PathMatch objects define patterns, contained inside PatternMatch objects (see Section 4.1.5), to match against the requested URI path. PatternMatch objects contain the pattern strings and flags that describe the URI path that a PathMatch applies to. PathMetadata objects contain the GenericMetadata objects that apply to content requests matching the defined URI path pattern. PathMetadata properties override properties previously defined in HostMetadata or less specific PathMatch paths. PathMetadata objects can contain additional PathMatch objects to recursively define more specific URI paths to which GenericMetadata properties might be applied.

A GenericMetadata object contains individual CDNI metadata objects which define the specific policies and attributes needed to properly deliver the associated content. For example, a GenericMetadata object could describe the source from which a CDN can acquire a piece of content. The GenericMetadata object is an atomic unit that can be referenced by HostMetadata or PathMetadata objects.

For example, if "example.com" is a content provider, a HostMatch object could include an entry for "example.com" with the URI of the associated HostMetadata object. The HostMetadata object for "example.com" describes the metadata properties which apply to "example.com" and could contain PathMatches for "example.com/movies/" and "example.com/music/", which in turn reference corresponding PathMetadata objects that contain the properties for those more specific URI paths. The PathMetadata object for "example.com/movies/" describes the properties which apply to that URI path. It could also contain a PathMatch object for

"example.com/movies/hd/*" which would reference the corresponding PathMetadata object for the "example.com/movies/hd/" path prefix.

The relationships in Figure 1 are also represented in tabular format in Table 1 below.

Data Object	Objects it contains or references
HostIndex	0 or more HostMatch objects.
HostMatch	1 HostMetadata object.
HostMetadata	0 or more PathMatch objects. 0 or more GenericMetadata objects.
PathMatch	1 PatternMatch object. 1 PathMetadata object.
PatternMatch	Does not contain or reference any other objects.
PathMetadata	0 or more PathMatch objects. 0 or more GenericMetadata objects.

Table 1: Relationships between CDNI Metadata Objects
(Table Representation)

3.2. Generic CDNI Metadata Objects

The HostMetadata and PathMetadata objects contain other CDNI metadata objects that contain properties which describe how User Agent requests for content should be processed, for example where to acquire the content from, authorization rules that should be applied, geo-blocking restrictions, and so on. Each such CDNI metadata object is a specialization of a CDNI GenericMetadata object. The GenericMetadata object abstracts the basic information required for metadata override and metadata distribution, from the specifics of any given property (i.e., property semantics, enforcement options, etc.).

The GenericMetadata object defines the properties contained within it as well as whether or not the properties are "mandatory-to-enforce". If the dCDN does not understand or support a "mandatory-to-enforce" property, the dCDN MUST NOT serve the content. If the property is not "mandatory-to-enforce", then that GenericMetadata object can be safely ignored and the content request can be processed in accordance with the rest of the CDNI metadata.

Although a CDN MUST NOT serve content to a User Agent if a "mandatory-to-enforce" property cannot be enforced, it could still be "safe-to-redistribute" that metadata to another CDN without modification. For example, in the cascaded CDN case, a transit CDN (tCDN) could pass through "mandatory-to-enforce" metadata to a dCDN.

For metadata which does not require customization or translation (i.e., metadata that is "safe-to-redistribute"), the data representation received off the wire MAY be stored and redistributed without being understood or supported by the transit CDN. However, for metadata which requires translation, transparent redistribution of the uCDN metadata values might not be appropriate. Certain metadata can be safely, though perhaps not optimally, redistributed unmodified. For example, source acquisition address might not be optimal if transparently redistributed, but it might still work.

Redistribution safety MUST be specified for each GenericMetadata property. If a CDN does not understand or support a given GenericMetadata property that is not "safe-to-redistribute", the CDN MUST set the "incomprehensible" flag to true for that GenericMetadata object before redistributing the metadata. The "incomprehensible" flag signals to a dCDN that the metadata was not properly transformed by the transit CDN. A CDN MUST NOT attempt to use metadata that has been marked as "incomprehensible" by a uCDN.

Transit CDNs MUST NOT change the value of "mandatory-to-enforce" or "safe-to-redistribute" when propagating metadata to a dCDN. Although a transit CDN can set the value of "incomprehensible" to true, a transit CDN MUST NOT change the value of "incomprehensible" from true to false.

Table 2 describes the action to be taken by a transit CDN (tCDN) for the different combinations of "mandatory-to-enforce" (MtE) and "safe-to-redistribute" (StR) properties, when the tCDN either does or does not understand the metadata in question:

MtE	StR	Metadata Understood by tCDN	Action
False	True	True	Can serve and redistribute.
False	True	False	Can serve and redistribute.
False	False	False	Can serve. MUST set "incomprehensible" to True when redistributing.
False	False	True	Can serve. Can redistribute after transforming the metadata (if the CDN knows how to do so safely), otherwise MUST set "incomprehensible" to True when redistributing.
True	True	True	Can serve and redistribute.
True	True	False	MUST NOT serve but can redistribute.
True	False	True	Can serve. Can redistribute after transforming the metadata (if the CDN knows how to do so safely), otherwise MUST set "incomprehensible" to True when redistributing.
True	False	False	MUST NOT serve. MUST set "incomprehensible" to True when redistributing.

Table 2: Action to be taken by a tCDN for the different combinations of MtE and StR properties

Table 3 describes the action to be taken by a dCDN for the different combinations of "mandatory-to-enforce" (MtE) and "incomprehensible" (Incomp) properties, when the dCDN either does or does not understand the metadata in question:

MtE	Incomp	Metadata Understood by dCDN	Action
False	False	True	Can serve.
False	True	True	Can serve but MUST NOT interpret/apply any metadata marked incomprehensible.
False	False	False	Can serve.
False	True	False	Can serve but MUST NOT interpret/apply any metadata marked incomprehensible.
True	False	True	Can serve.
True	True	True	MUST NOT serve.
True	False	False	MUST NOT serve.
True	True	False	MUST NOT serve.

Table 3: Action to be taken by a dCDN for the different combinations of MtE and Incomp properties

3.3. Metadata Inheritance and Override

In the metadata object model, a HostMetadata object can contain multiple PathMetadata objects (via PathMatch objects). Each PathMetadata object can in turn contain other PathMetadata objects. HostMetadata and PathMetadata objects form an inheritance tree where each node in the tree inherits or overrides the property values set by its parent.

GenericMetadata objects of a given type override all GenericMetadata objects of the same type previously defined by any parent object in the tree. GenericMetadata objects of a given type previously defined by a parent object in the tree are inherited when no object of the same type is defined by the child object. For example, if HostMetadata for the host "example.com" contains GenericMetadata objects of type LocationACL and TimeWindowACL, while a PathMetadata object which applies to "example.com/movies/*" defines an alternate GenericMetadata object of type TimeWindowACL, then:

- o the TimeWindowACL defined in the PathMetadata would override the TimeWindowACL defined in the HostMetadata for all User Agent requests for content under "example.com/movies/", and
- o the LocationACL defined in the HostMetadata would be inherited for all User Agent requests for content under "example.com/movies/".

A single HostMetadata or PathMetadata object MUST NOT contain multiple GenericMetadata objects of the same type. If an array of GenericMetadata contains objects of duplicate types, the receiver MUST ignore all but the first object of each type.

4. CDNI Metadata objects

Section 4.1 provides the definitions of each metadata object type introduced in Section 3. These metadata objects are described as structural metadata objects as they provide the structure for host and URI path-based inheritance and identify which GenericMetadata objects apply to a given User Agent content request.

Section 4.2 provides the definitions for a base set of core metadata objects which can be contained within a GenericMetadata object. These metadata objects govern how User Agent requests for content are handled. GenericMetadata objects can contain other GenericMetadata as properties; these can be referred to as sub-objects). As with all CDNI metadata objects, the value of the GenericMetadata sub-objects can be either a complete serialized representation of the sub-object, or a Link object that contains a URI that can be dereferenced to retrieve the complete serialized representation of the property sub-object.

Section 6.5 discusses the ability to extend the base set of GenericMetadata objects specified in this document with additional standards-based or vendor specific GenericMetadata objects that might be defined in the future in separate documents.

dCDNs and tCDNs MUST support parsing of all CDNI metadata objects specified in this document. A dCDN does not have to implement the underlying functionality represented by non-structural GenericMetadata objects (though that might restrict the content that a given dCDN will be able to serve). uCDNs as generators of CDNI metadata only need to support generating the CDNI metadata that they need in order to express the policies required by the content they are describing. See Section 6.4 for more details on the specific encoding rules for CDNI metadata objects.

Note: In the following sections, the term "mandatory-to-specify" is used to convey which properties MUST be included for a given structural or GenericMetadata object. When mandatory-to-specify is specified as "Yes" for an individual property, it means that if the object containing that property is included in a metadata response, then the mandatory-to-specify property MUST also be included (directly or by reference) in the response, e.g., a HostMatch property object without a host to match against does not make sense,

therefore, the host property is mandatory-to-specify inside a HostMatch object.

4.1. Definitions of the CDNI structural metadata objects

Each of the sub-sections below describe the structural objects introduced in Section 3.1.

4.1.1. HostIndex

The HostIndex object is the entry point into the CDNI metadata hierarchy. It contains an array of HostMatch objects. An incoming content request is checked against the Hostname (or IP address) specified by each of the listed HostMatch objects to find the HostMatch object which applies to the request.

Property: hosts

Description: Array of HostMatch objects. Hosts (HostMatch objects) MUST be evaluated in the order they appear and the first HostMatch object that matches the content request being processed MUST be used.

Type: Array of HostMatch objects

Mandatory-to-Specify: Yes.

Example HostIndex object containing two HostMatch objects, where the first HostMatch object is embedded and the second HostMatch object is referenced:

```
{
  "hosts": [
    {
      <Properties of embedded HostMatch object>
    },
    {
      "type": "MI.HostMatch",
      "href": "https://metadata.ucdn.example/hostmatch1234"
    }
  ]
}
```

4.1.2. HostMatch

The HostMatch object contains a Hostname or IP address to match against content requests. The HostMatch object also contains a HostMetadata object to apply if a match is found.

Property: host

Description: Hostname or IP address and optional port to match against the requested host, i.e., the [RFC3986] host and port. In order for a Hostname or IP address in a content request to match the Hostname or IP address in the host property the value from the content request when converted to lowercase MUST be identical to the value of the host property when converted to lowercase. All implementations MUST support IPv4 addresses encoded as specified by the 'IPv4address' rule in Section 3.2.2 of [RFC3986]. IPv6 addresses MUST be encoded in one of the IPv6 address formats specified in [RFC5952] although receivers MUST support all IPv6 address formats specified in [RFC4291]. Hostnames MUST conform to the Domain Name System (DNS) syntax defined in [RFC1034] and [RFC1123]. Internationalized Domain Names (IDN) must first be transformed to the the A-label form [RFC5890] as per [RFC5891].

Type: Endpoint

Mandatory-to-Specify: Yes.

Property: host-metadata

Description: CDNI metadata to apply when delivering content that matches this host.

Type: HostMetadata

Mandatory-to-Specify: Yes.

Example HostMatch object with an embedded HostMetadata object:

```
{
  "host": "video.example.com",
  "host-metadata" : {
    <Properties of embedded HostMetadata object>
  }
}
```

Example HostMatch object referencing (via a Link object, see Section 4.3.1) a HostMetadata object:

```
{
  "host": "video.example.com",
  "host-metadata" : {
    "type": "MI.HostMetadata",
    "href": "https://metadata.ucdn.example/host1234"
  }
}
```

4.1.3. HostMetadata

A HostMetadata object contains the CDNI metadata properties for content served for a particular host (defined in the HostMatch object) and possibly child PathMatch objects.

Property: metadata

Description: Array of host related metadata.

Type: Array of GenericMetadata objects

Mandatory-to-Specify: Yes.

Property: paths

Description: Path specific rules. Path patterns (PathMatch objects) MUST be evaluated in the order they appear and the first (and only the first) PathMatch object that matches the content request being processed MUST be used.

Type: Array of PathMatch objects

Mandatory-to-Specify: No.

Example HostMetadata object containing a number of embedded GenericMetadata objects that will describe the default metadata for the host and an embedded PathMatch object that contains a path for which metadata exists that overrides the default metadata for the host:


```
{
  "metadata": [
    {
      <Properties of 1st embedded GenericMetadata object>
    },
    {
      <Properties of 2nd embedded GenericMetadata object>
    },
    ...
    {
      <Properties of Nth embedded GenericMetadata object>
    }
  ],
  "paths": [
    {
      <Properties of embedded PathMatch object>
    }
  ]
}
```

4.1.4. PathMatch

A PathMatch object contains PatternMatch object with a path to match against a resource's URI path, as well as how to handle URI query parameters. The PathMatch also contains a PathMetadata object with GenericMetadata to apply if the resource's URI matches the pattern within the PatternMatch object.

Property: path-pattern

Description: Pattern to match against the requested resource's URI.

Type: PatternMatch

Mandatory-to-Specify: Yes.

Property: path-metadata

Description: CDNI metadata to apply when delivering content that matches the associated PatternMatch.

Type: PathMetadata

Mandatory-to-Specify: Yes.

Example PathMatch object referencing the PathMetadata object to use for URIs that match the case-sensitive URI path pattern `"/movies/*"` (contained within an embedded PatternMatch object):

```
{
  "path-pattern": {
    "pattern": "/movies/*",
    "case-sensitive": true
  },
  "path-metadata": {
    "type": "MI.PathMetadata",
    "href": "https://metadata.ucdn.example/host1234/pathDCE"
  }
}
```

4.1.5. PatternMatch

A PatternMatch object contains the pattern string and flags that describe the pattern expression.

Property: pattern

Description: A pattern for matching against the URI path, i.e., against the [RFC3986] path-absolute. The pattern can contain the wildcards `*` and `?`, where `*` matches any sequence of [RFC3986] pchar or `"/` characters (including the empty string) and `?` matches exactly one [RFC3986] pchar character. The three literals `$`, `*` and `?` MUST be escaped as `$$`, `$*` and `$?` (where `$` is the designated escape character). All other characters are treated as literals.

Type: String

Mandatory-to-Specify: Yes.

Property: case-sensitive

Description: Flag indicating whether or not case-sensitive matching should be used. Note: Case-insensitivity applies to ALPHA characters in the URI path prior to percent-decoding [RFC3986].

Type: Boolean

Mandatory-to-Specify: No. Default is case-insensitive match.

Example PatternMatch object that matches the case-sensitive URI path pattern `"/movies/*"`. All query parameters will be ignored when

matching URIs requested from surrogates by content clients against this path pattern:

```
{
  "pattern": "/movies/*",
  "case-sensitive": true
}
```

Example PatternMatch object that matches the case-sensitive URI path pattern "/movies/*". Only the query parameter "sessionid" will be evaluated when matching URIs requested from surrogates by content clients against this path pattern:

```
{
  "pattern": "/movies/*",
  "case-sensitive": true
}
```

4.1.6. PathMetadata

A PathMetadata object contains the CDNI metadata properties for content requests that match against the associated URI path (defined in a PathMatch object).

Note that if DNS-based redirection is employed, then a dCDN will be unable to evaluate any metadata at the PathMetadata level or below because only the hostname of the content request is available at request routing time. dCDNs SHOULD still process all PathMetadata for the host before responding to the redirection request to detect if any unsupported metadata is specified. If any metadata not supported by the dCDN is marked as "mandatory-to-enforce", the dCDN SHOULD NOT accept the content redirection request, in order to avoid receiving content requests that it will not be able to satisfy/serve.

Property: metadata

Description: Array of path related metadata.

Type: Array of GenericMetadata objects

Mandatory-to-Specify: Yes.

Property: paths

Description: Path specific rules. Path patterns (PathMatch objects) MUST be evaluated in the order they appear and the first (and only the first) PathMatch object that matches the content request being processed MUST be used.

Type: Array of PathMatch objects

Mandatory-to-Specify: No.

Example PathMetadata object containing a number of embedded GenericMetadata objects that describe the metadata to apply for the URI path defined in the parent PathMatch object, as well as a more specific PathMatch object.

```
{
  "metadata": [
    {
      <Properties of 1st embedded GenericMetadata object>
    },
    {
      <Properties of 2nd embedded GenericMetadata object>
    },
    ...
    {
      <Properties of Nth embedded GenericMetadata object>
    }
  ],
  "paths": [
    {
      <Properties of embedded PathMatch object>
    }
  ]
}
```

4.1.7. GenericMetadata

A GenericMetadata object is a wrapper for managing individual CDNI metadata properties in an opaque manner.

Property: generic-metadata-type

Description: Case-insensitive CDNI metadata object type.

Type: String containing the CDNI Payload Type [RFC7736] of the object contained in the generic-metadata-value property (see Table 4).

Mandatory-to-Specify: Yes.

Property: generic-metadata-value

Description: CDNI metadata object.

Type: Format/Type is defined by the value of generic-metadata-type property above. Note: generic-metadata-values MUST NOT name any properties "href" (see Section 4.3.1).

Mandatory-to-Specify: Yes.

Property: mandatory-to-enforce

Description: Flag identifying whether or not the enforcement of the property metadata is required.

Type: Boolean

Mandatory-to-Specify: No. Default is to treat metadata as mandatory to enforce (i.e., a value of True).

Property: safe-to-redistribute

Description: Flag identifying whether or not the property metadata can be safely redistributed without modification.

Type: Boolean

Mandatory-to-Specify: No. Default is allow transparent redistribution (i.e., a value of True).

Property: incomprehensible

Description: Flag identifying whether or not any CDN in the chain of delegation has failed to understand and/or failed to properly transform this metadata object. Note: This flag only applies to metadata objects whose safe-to-redistribute property has a value of False.

Type: Boolean

Mandatory-to-Specify: No. Default is comprehensible (i.e., a value of False).

Example GenericMetadata object containing a metadata object that applies to the applicable URI path and/or host (within a parent PathMetadata and/or HostMetadata object, respectively):

```
{
  "mandatory-to-enforce": true,
  "safe-to-redistribute": true,
  "incomprehensible": false,
  "generic-metadata-type": <CDNI Payload Type of this metadata object>,
  "generic-metadata-value":
    {
      <Properties of this metadata object>
    }
}
```

4.2. Definitions of the initial set of CDNI Generic Metadata objects

The objects defined below are intended to be used in the GenericMetadata object generic-metadata-value field as defined in Section 4.1.7 and their generic-metadata-type property MUST be set to the appropriate CDNI Payload Type as defined in Table 4.

4.2.1. SourceMetadata

Source metadata provides the dCDN with information about content acquisition, i.e., how to contact an uCDN Surrogate or an Origin Server to obtain the content to be served. The sources are not necessarily the actual Origin Servers operated by the CSP but might be a set of Surrogates in the uCDN.

Property: sources

Description: Sources from which the dCDN can acquire content, listed in order of preference.

Type: Array of Source objects (see Section 4.2.1.1)

Mandatory-to-Specify: No. Default is to use static configuration, out-of-band from the metadata interface.

Example SourceMetadata object (which contains two Source objects) that describes which servers the dCDN should use for acquiring content for the applicable URI path and/or host:

```
{
  "generic-metadata-type": "MI.SourceMetadata",
  "generic-metadata-value":
    {
      "sources": [
        {
          "endpoints": [
            "a.servicel23.ucdn.example",
            "b.servicel23.ucdn.example"
          ],
          "protocol": "http/1.1"
        },
        {
          "endpoints": ["origin.servicel23.example"],
          "protocol": "http/1.1"
        }
      ]
    }
}
```

4.2.1.1. Source

A Source object describes the source to be used by the dCDN for content acquisition (e.g., a Surrogate within the uCDN or an alternate Origin Server), the protocol to be used, and any authentication method to be used when contacting that source.

Endpoints within a Source object **MUST** be treated as equivalent/equal. A uCDN can specify an array of sources in preference order within a SourceMetadata object, and then for each preference ranked Source object, a uCDN can specify an array of endpoints that are equivalent (e.g., a pool of servers that are not behind a load balancer).

Property: acquisition-auth

Description: Authentication method to use when requesting content from this source.

Type: Auth (see Section 4.2.7)

Mandatory-to-Specify: No. Default is no authentication required.

Property: endpoints

Description: Origins from which the dCDN can acquire content. If multiple endpoints are specified they are all equal, i.e.,

the list is not in preference order (e.g., a pool of servers behind a load balancer).

Type: Array of Endpoint objects (See Section 4.3.3)

Mandatory-to-Specify: Yes.

Property: protocol

Description: Network retrieval protocol to use when requesting content from this source.

Type: Protocol (see Section 4.3.2)

Mandatory-to-Specify: Yes.

Example Source object that describes a pair of endpoints (servers) the dCDN can use for acquiring content for the applicable host and/or URI path:

```
{
  "endpoints": [
    "a.servicel23.ucdn.example",
    "b.servicel23.ucdn.example"
  ],
  "protocol": "http/1.1"
}
```

4.2.2.2. LocationACL Metadata

LocationACL metadata defines which locations a User Agent needs to be in, in order to be able to receive the associated content.

A LocationACL which does not include a locations property results in an action of allow all, meaning that delivery can be performed regardless of the User Agent's location, otherwise a CDN MUST take the action from the first footprint to match against the User Agent's location. If two or more footprints overlap, the first footprint that matches against the User Agent's location determines the action a CDN MUST take. If the locations property is included but is empty, or if none of the listed footprints matches the User Agent's location, then the result is an action of deny.

Although the LocationACL, TimeWindowACL (see Section 4.2.3), and ProtocolACL (see Section 4.2.4) are independent GenericMetadata objects, they can provide conflicting information to a dCDN, e.g., a content request which is simultaneously allowed based on the LocationACL and denied based on the TimeWindowACL. The dCDN MUST use

the logical AND of all ACLs (where 'allow' is true and 'deny' is false) to determine whether or not a request should be allowed.

Property: locations

Description: Access control list which allows or denies (blocks) delivery based on the User Agent's location.

Type: Array of LocationRule objects (see Section 4.2.2.1)

Mandatory-to-Specify: No. Default is allow all locations.

Example LocationACL object that allows the dCDN to deliver content to any location/IP address:

```
{
  "generic-metadata-type": "MI.LocationACL",
  "generic-metadata-value":
    {
    }
}
```

Example LocationACL object (which contains a LocationRule object which itself contains a Footprint object) that only allows the dCDN to deliver content to User Agents in the USA:

```
{
  "generic-metadata-type": "MI.LocationACL",
  "generic-metadata-value":
    {
      "locations": [
        {
          "action": "allow",
          "footprints": [
            {
              "footprint-type": "countrycode",
              "footprint-value": ["us"]
            }
          ]
        }
      ]
    }
}
```

4.2.2.1. LocationRule

A LocationRule contains or references an array of Footprint objects and the corresponding action.

Property: footprints

Description: Array of footprints to which the rule applies.

Type: Array of Footprint objects (see Section 4.2.2.2)

Mandatory-to-Specify: Yes.

Property: action

Description: Defines whether the rule specifies locations to allow or deny.

Type: Enumeration [allow|deny] encoded as a lowercase string

Mandatory-to-Specify: No. Default is deny.

Example LocationRule object (which contains a Footprint object) that allows the dCDN to deliver content to clients in the USA:

```
{
  "action": "allow",
  "footprints": [
    {
      "footprint-type": "countrycode",
      "footprint-value": ["us"]
    }
  ]
}
```

4.2.2.2. Footprint

A Footprint object describes the footprint to which a LocationRule can be applied to, e.g., an IPv4 address range or a geographic location.

Property: footprint-type

Description: Registered footprint type (see Section 7.2). The footprint types specified by this document are: "ipv4cidr" (IPv4CIDR, see Section 4.3.5), "ipv6cidr" (IPv6CIDR, see Section 4.3.6), "asn" (Autonomous System Number, see

Section 4.3.7) and "countrycode" (Country Code, see Section 4.3.8).

Type: Lowercase String

Mandatory-to-Specify: Yes.

Property: footprint-value

Description: Array of footprint values conforming to the specification associated with the registered footprint type. Footprint values can be simple strings (e.g., IPv4CIDR, IPv6CIDR, ASN, and CountryCode), however, other Footprint objects can be defined in the future, along with a more complex encoding (e.g., GPS coordinate tuples).

Type: Array of footprints

Mandatory-to-Specify: Yes.

Example Footprint object describing a footprint covering the USA:

```
{
  "footprint-type": "countrycode",
  "footprint-value": ["us"]
}
```

Example Footprint object describing a footprint covering the IP address ranges 192.0.2.0/24 and 198.51.100.0/24:

```
{
  "footprint-type": "ipv4cidr",
  "footprint-value": ["192.0.2.0/24", "198.51.100.0/24"]
}
```

Example Footprint object describing a footprint covering the IP address ranges 2001:db8::/32:

```
{
  "footprint-type": "ipv6cidr",
  "footprint-value": ["2001:db8::/32"]
}
```

Example Footprint object describing a footprint covering the autonomous system 64496:

```
{
  "footprint-type": "asn",
  "footprint-value": ["as64496"]
}
```

4.2.3. TimeWindowACL

TimeWindowACL metadata defines time-based restrictions.

A TimeWindowACL which does not include a times property results in an action of allow all, meaning that delivery can be performed regardless of the time of the User Agent's request, otherwise a CDN MUST take the action from the first window to match against the current time. If two or more windows overlap, the first window that matches against the current time determines the action a CDN MUST take. If the times property is included but is empty, or if none of the listed windows matches the current time, then the result is an action of deny.

Although the LocationACL (see Section 4.2.2), TimeWindowACL, and ProtocolACL (see Section 4.2.4) are independent GenericMetadata objects, they can provide conflicting information to a dCDN, e.g., a content request which is simultaneously allowed based on the LocationACL and denied based on the TimeWindowACL. The dCDN MUST use the logical AND of all ACLs (where 'allow' is true and 'deny' is false) to determine whether or not a request should be allowed.

Property: times

Description: Access control list which allows or denies (blocks) delivery based on the time of a User Agent's request.

Type: Array of TimeWindowRule objects (see Section 4.2.3.1)

Mandatory-to-Specify: No. Default is allow all time windows.

Example TimeWindowACL object (which contains a TimeWindowRule object which itself contains a TimeWindow object) that only allows the dCDN to deliver content to clients between 09:00 01/01/2000 UTC and 17:00 01/01/2000 UTC:

```
{
  "generic-metadata-type": "MI.TimeWindowACL",
  "generic-metadata-value":
    {
      "times": [
        {
          "action": "allow",
          "windows": [
            {
              "start": 946717200,
              "end": 946746000
            }
          ]
        }
      ]
    }
}
```

4.2.3.1. TimeWindowRule

A TimeWindowRule contains or references an array of TimeWindow objects and the corresponding action.

Property: windows

Description: Array of time windows to which the rule applies.

Type: Array of TimeWindow objects (see Section 4.2.3.2)

Mandatory-to-Specify: Yes.

Property: action

Description: Defines whether the rule specifies time windows to allow or deny.

Type: Enumeration [allow|deny] encoded as a lowercase string

Mandatory-to-Specify: No. Default is deny.

Example TimeWindowRule object (which contains a TimeWindow object) that only allows the dCDN to deliver content to clients between 09:00 01/01/2000 UTC and 17:00 01/01/2000 UTC:

```
{
  "action": "allow",
  "windows": [
    {
      "start": 946717200,
      "end": 946746000
    }
  ]
}
```

4.2.3.2. TimeWindow

A TimeWindow object describes a time range which can be applied by an TimeWindowACL, e.g., start 946717200 (i.e., 09:00 01/01/2000 UTC), end: 946746000 (i.e., 17:00 01/01/2000 UTC).

Property: start

Description: The start time of the window.

Type: Time (see Section 4.3.4)

Mandatory-to-Specify: Yes.

Property: end

Description: The end time of the window.

Type: Time (see Section 4.3.4)

Mandatory-to-Specify: Yes.

Example TimeWindow object that describes a time window from 09:00 01/01/2000 UTC to 17:00 01/01/2000 UTC:

```
{
  "start": 946717200,
  "end": 946746000
}
```

4.2.4. ProtocolACL Metadata

ProtocolACL metadata defines delivery protocol restrictions.

A ProtocolACL which does not include a protocol-acl property results in an action of allow all, meaning that delivery can be performed regardless of the protocol in the User Agent's request, otherwise a CDN MUST take the action from the first protocol to match against the

request protocol. If two or more request protocols overlap, the first protocol that matches the request protocol determines the action a CDN MUST take. If the protocol-acl property is included but is empty, or if none of the listed protocol matches the request protocol, then the result is an action of deny.

Although the LocationACL, TimeWindowACL, and ProtocolACL are independent GenericMetadata objects, they can provide conflicting information to a dCDN, e.g., a content request which is simultaneously allowed based on the ProtocolACL and denied based on the TimeWindowACL. The dCDN MUST use the logical AND of all ACLs (where 'allow' is true and 'deny' is false) to determine whether or not a request should be allowed.

Property: protocol-acl

Description: Description: Access control list which allows or denies (blocks) delivery based on delivery protocol.

Type: Array of ProtocolRule objects (see Section 4.2.4.1)

Mandatory-to-Specify: No. Default is allow all protocols.

Example ProtocolACL object (which contains a ProtocolRule object) that only allows the dCDN to deliver content using HTTP/1.1:

```
{
  "generic-metadata-type": "MI.ProtocolACL",
  "generic-metadata-value":
    {
      "protocol-acl": [
        {
          "action": "allow",
          "protocols": ["http/1.1"]
        }
      ]
    }
}
```

4.2.4.1. ProtocolRule

A ProtocolRule contains or references an array of Protocol objects and the corresponding action.

Property: protocols

Description: Array of protocols to which the rule applies.

Type: Array of Protocols (see Section 4.3.2)

Mandatory-to-Specify: Yes.

Property: action

Description: Defines whether the rule specifies protocols to allow or deny.

Type: Enumeration [allow|deny] encoded as a lowercase string

Mandatory-to-Specify: No. Default is deny.

Example ProtocolRule object (which contains a ProtocolRule object) that allows the dCDN to deliver content using HTTP/1.1:

```
{
  "action": "allow",
  "protocols": ["http/1.1"]
}
```

4.2.5. DeliveryAuthorization Metadata

Delivery Authorization defines authorization methods for the delivery of content to User Agents.

Property: delivery-auth-methods

Description: Options for authorizing content requests. Delivery for a content request is authorized if any of the authorization methods in the list is satisfied for that request.

Type: Array of Auth objects (see Section 4.2.7)

Mandatory-to-Specify: No. Default is no authorization required.

Example DeliveryAuthorization object (which contains an Auth object):


```

{
  "generic-metadata-type": "MI.DeliveryAuthorization",
  "generic-metadata-value":
  {
    "delivery-auth-methods": [
      {
        "auth-type": <CDNI Payload Type of this Auth object>,
        "auth-value":
        {
          <Properties of this Auth object>
        }
      }
    ]
  }
}

```

4.2.6. Cache

A Cache object describes the cache control parameters to be applied to the content by intermediate caches.

Cache keys are generated from the URI of the content request [RFC7234]. In some cases, a CDN or content provider might want certain path segments or query parameters to be excluded from the cache key generation. The Cache object provides guidance on what parts of the path and query string to include.

Property: exclude-path-pattern

Description: A pattern for matching against the URI path, i.e., against the [RFC3986] path-absolute. The pattern can contain the wildcards * and ?, where * matches any sequence of [RFC3986] pchar or "/" characters (including the empty string) and ? matches exactly one [RFC3986] pchar character. The three literals \$, * and ? MUST be escaped as \$\$, \$* and \$? (where \$ is the designated escape character). All other characters are treated as literals. Cache key generation MUST only include the portion of the path-absolute that matches the wildcard portions of the pattern. Note: Inconsistency between the PatternMatch pattern Section 4.1.5 and the exclude-path-pattern can result in inefficient caching.

Type: String

Mandatory-to-Specify: No. Default is to use the full URI path-absolute to generate the cache key.

Property: include-query-strings

Description: Allows a Surrogate to specify the URI query string parameters [RFC3986] to include when comparing the requested URI against the URIs in its cache for equivalence. Matching query parameters MUST be case-insensitive. If all query parameters should be ignored, then the list MUST be specified and MUST be empty. If a query parameter appears multiple times in the query string, each instance value MUST be aggregated prior to comparison. For consistent cache key generation, query parameters SHOULD be evaluated in the order specified in this array.

Type: Array of String

Mandatory-to-Specify: No. Default is to consider all query string parameters when comparing URIs.

Example Cache object that instructs the dCDN to use the full URI path and ignore all query parameters:

```
{
  "generic-metadata-type": "MI.Cache",
  "generic-metadata-value":
  {
    "include-query-strings": []
  }
}
```

Example Cache object that instructs the dCDN to exclude the "CDNX" path prefix and only include the (case-insensitive) query parameters named "mediaid" and "providerid":

```
{
  "generic-metadata-type": "MI.Cache",
  "generic-metadata-value":
  {
    "exclude-path-pattern": "/CDNX/*",
    "include-query-strings": ["mediaid", "providerid"]
  }
}
```

Example Cache object that instructs the dCDN to exclude the "CDNX" path prefix, but includes all query parameters:

```
{
  "generic-metadata-type": "MI.Cache",
  "generic-metadata-value":
  {
    "exclude-path-pattern": "/CDNX/*"
  }
}
```

4.2.7. Auth

An Auth object defines authentication and authorization methods to be used during content acquisition and content delivery, respectively.

Note: This document does not define any Auth methods. Individual Auth methods are being defined separately (e.g., URI Signing [I-D.ietf-cdni-uri-signing]). The GenericMetadata which contain Auth objects is defined herein for convenience and so as not to be specific to any particular Auth method.

Property: auth-type

Description: Auth type (The CDNI Payload Type [RFC7736] of the GenericMetadata object contained in the auth-value property).

Type: String

Mandatory-to-Specify: Yes.

Property: auth-value

Description: An object conforming to the specification associated with the Auth type.

Type: GenericMetadata Object

Mandatory-to-Specify: Yes.

Example Auth object:

```
{
  "generic-metadata-type": "MI.Auth",
  "generic-metadata-value":
  {
    "auth-type": <CDNI Payload Type of this Auth object>,
    "auth-value":
    {
      <Properties of this Auth object>
    }
  }
}
```

4.2.8. Grouping

A Grouping object identifies a group of content to which a given asset belongs.

Property: ccid

Description: Content Collection identifier for an application-specific purpose such as logging aggregation.

Type: String

Mandatory-to-Specify: No. Default is an empty string.

Example Grouping object that specifies a Content Collection Identifier for the content associated with the Grouping object's parent HostMetadata and PathMetadata:

```
{
  "generic-metadata-type": "MI.Grouping",
  "generic-metadata-value":
  {
    "ccid": "ABCD"
  }
}
```

4.3. CDNI Metadata Simple Data Type Descriptions

This section describes the simple data types that are used for properties of CDNI metadata objects.

4.3.1. Link

A Link object can be used in place of any of the objects or properties described above. Link objects can be used to avoid duplication if the same metadata information is repeated within the

metadata tree. When a Link object replaces another object, its href property is set to the URI of the resource and its type property is set to the CDNI Payload Type of the object it is replacing.

dCDNs can detect the presence of a Link object by detecting the presence of a property named "href" within the object. This means that GenericMetadata types MUST NOT contain a property named "href" because doing so would conflict with the ability for dCDNs to detect Link objects being used to reference a GenericMetadata object.

Property: href

Description: The URI of the addressable object being referenced.

Type: String

Mandatory-to-Specify: Yes.

Property: type

Description: The CDNI Payload type of the object being referenced.

Type: String

Mandatory-to-Specify: No. If the container specifies the type (e.g., the HostIndex object contains an array of HostMatch objects, so a Link object in the list of HostMatch objects must reference a HostMatch), then it is not necessary to explicitly specify a type.

Example Link object referencing a HostMatch object:

```
{
  "type": "MI.HostMatch",
  "href": "https://metadata.ucdn.example/hostmatch1234"
}
```

Example Link object referencing a HostMatch object, without an explicit type, inside a HostIndex object:

```
{
  "hosts": [
    {
      <Properties of embedded HostMatch object>
    },
    {
      "href": "https://metadata.ucdn.example/hostmatch1234"
    }
  ]
}
```

4.3.1.1. Link Loop Prevention

When following a Link, CDNI metadata clients SHOULD verify that the CDNI Payload Type of the object retrieved matches the expected CDNI Payload Type, as indicated by the link object. For GenericMetadata objects, type checks will prevent self references; however, incorrect linking can result in circular references for structural metadata objects, specifically, PathMatch and PathMetadata objects Figure 1. To prevent the circular references, CDNI metadata clients SHOULD verify that no duplicate Links occur for PathMatch or PathMetadata objects.

4.3.2. Protocol

Protocol objects are used to specify registered protocols for content acquisition or delivery (see Section 7.3).

Type: String

Example:

"http/1.1"

4.3.3. Endpoint

A Hostname (with optional port) or an IP address (with optional port).

All implementations MUST support IPv4 addresses encoded as specified by the 'IPv4address' rule in Section 3.2.2 of [RFC3986]. IPv6 addresses MUST be encoded in one of the IPv6 address formats specified in [RFC5952] although receivers MUST support all IPv6 address formats specified in [RFC4291]. Hostnames MUST conform to the Domain Name System (DNS) syntax defined in [RFC1034] and [RFC1123]. Internationalized Domain Names (IDN) must first be transformed to the A-label form [RFC5890] as per [RFC5891].

Type: String

Example Hostname:

"metadata.ucdn.example"

Example IPv4 address:

"192.0.2.1"

Example IPv6 address (with port number):

"[2001:db8::1]:81"

4.3.4. Time

A time value expressed in seconds since the Unix epoch (i.e., zero hours, zero minutes, zero seconds, on January 1, 1970) Coordinated Universal Time (UTC) [POSIX].

Type: Integer

Example Time representing 09:00:00 01/01/2000 UTC:

946717200

4.3.5. IPv4CIDR

An IPv4address CIDR block encoded as specified by the 'IPv4address' rule in Section 3.2.2 of [RFC3986] followed by a / followed by an unsigned integer representing the leading bits of the routing prefix (i.e., IPv4 CIDR notation). Single IP addresses can be expressed as /32.

Type: String

Example IPv4 CIDR:

"192.0.2.0/24"

4.3.6. IPv6CIDR

An IPv6address CIDR block encoded in one of the IPv6 address formats specified in [RFC5952] followed by a / followed by an unsigned integer representing the leading bits of the routing prefix (i.e., IPv6 CIDR notation). Single IP addresses can be expressed as /128.

Type: String

Example IPv6 CIDR:

"2001:db8::/32"

4.3.7. ASN

An Autonomous System Number encoded as a string consisting of the characters "as" (in lowercase) followed by the Autonomous System number [RFC6793].

Type: String

Example ASN:

"as64496"

4.3.8. CountryCode

An ISO 3166-1 alpha-2 code [ISO3166-1] in lowercase.

Type: String

Example Country Code representing the USA:

"us"

5. CDNI Metadata Capabilities

CDNI metadata is used to convey information pertaining to content delivery from uCDN to dCDN. For optional metadata, it can be useful for the uCDN to know if the dCDN supports the underlying functionality described by the metadata, prior to delegating any content requests to the dCDN. If some metadata is "mandatory-to-enforce", and the dCDN does not support it, any delegated requests for content that requires that metadata will fail. The uCDN will likely want to avoid delegating those requests to that dCDN. Likewise, for any metadata which might be assigned optional values, it could be useful for the uCDN to know which values a dCDN supports, prior to delegating any content requests to that dCDN. If the optional value assigned to a given piece of content's metadata is not supported by the dCDN, any delegated requests for that content can fail, so again the uCDN is likely to want to avoid delegating those requests to that dCDN.

The CDNI Footprint and Capabilities Interface (FCI) provides a means of advertising capabilities from dCDN to uCDN [RFC7336]. Support for optional metadata types and values can be advertised using the FCI.

6. CDNI Metadata interface

This section specifies an interface to enable a dCDN to retrieve CDNI metadata objects from a uCDN.

The interface can be used by a dCDN to retrieve CDNI metadata objects either:

- o Dynamically as required by the dCDN to process received requests. For example in response to a query from an uCDN over the CDNI Request Routing Redirection interface (RI) [I-D.ietf-cdni-redirection] or in response to receiving a request for content from a User Agent. Or;
- o In advance of being required. For example in the case of pre-positioned CDNI metadata acquisition, initiated through the "CDNI Control interface / Triggers" (CI/T) interface [I-D.ietf-cdni-control-triggers].

The CDNI metadata interface is built on the principles of HTTP web services. In particular, this means that requests and responses over the interface are built around the transfer of representations of hyperlinked resources. A resource in the context of the CDNI metadata interface is any object in the object model (as described in Section 3 and Section 4).

To retrieve CDNI metadata, a CDNI metadata client (i.e., a client in the dCDN) first makes a HTTP GET request for the URI of the HostIndex which provides the CDNI metadata client with an array of Hostnames for which the uCDN can delegate content delivery to the dCDN. The CDNI metadata client can then obtain any other CDNI metadata objects by making a HTTP GET requests for any linked metadata objects it requires.

CDNI metadata servers (i.e., servers in the uCDN) are free to assign whatever structure they desire to the URIs for CDNI metadata objects and CDNI metadata clients MUST NOT make any assumptions regarding the structure of CDNI metadata URIs or the mapping between CDNI metadata objects and their associated URIs. Therefore any URIs present in the examples in this document are purely illustrative and are not intended to impose a definitive structure on CDNI metadata interface implementations.

6.1. Transport

The CDNI metadata interface uses HTTP as the underlying protocol transport [RFC7230].

The HTTP Method in the request defines the operation the request would like to perform. A server implementation of the CDNI metadata interface MUST support the HTTP GET and HEAD methods.

The corresponding HTTP Response returns the status of the operation in the HTTP Status Code and returns the current representation of the resource (if appropriate) in the Response Body. HTTP Responses that contain a response body SHOULD include an ETag to enable validation of cached versions of returned resources.

As the CDNI metadata interface builds on top of HTTP, CDNI metadata server implementations MAY make use of any HTTP feature when implementing the CDNI metadata interface, for example, a CDNI metadata server MAY make use of HTTP's caching mechanisms to indicate that the returned response/representation can be reused without re-contacting the CDNI metadata server.

6.2. Retrieval of CDNI Metadata resources

In the general case, a CDNI metadata server makes CDNI metadata objects available via a unique URIs and thus, in order to retrieve CDNI metadata, a CDNI metadata client first makes a HTTP GET request for the URI of the HostIndex which provides an array of Hostnames for which the uCDN can delegate content delivery to the dCDN.

In order to retrieve the CDNI metadata for a particular request the CDNI metadata client processes the received HostIndex object and finds the corresponding HostMetadata entry (by matching the hostname in the request against the hostnames listed in the HostMatch objects). If the HostMetadata is linked (rather than embedded), the CDNI metadata client then makes a GET request for the URI specified in the href property of the Link object which points to the HostMetadata object itself.

In order to retrieve the most specific metadata for a particular request, the CDNI metadata client inspects the HostMetadata for references to more specific PathMetadata objects (by matching the URI path in the request against the path-patterns in any PathMatch objects listed in the HostMetadata object). If any PathMetadata are found to match (and are linked rather than embedded), the CDNI metadata client makes another GET request for the PathMetadata. Each PathMetadata object can also include references to yet more specific metadata. If this is the case, the CDNI metadata client continues requesting PathMatch and PathMetadata objects recursively. The CDNI metadata client repeats this approach of processing metadata objects and retrieving (via HTTP GETs) any linked objects until it has all the metadata objects it requires in order to process the redirection request from an uCDN or the content request from a User Agent.

In cases where a dCDN is not able to retrieve the entire set of CDNI metadata associated with a User Agent request, or it has retrieved that metadata but it is stale according to standard HTTP caching rules and cannot be revalidated, for example because the uCDN is unreachable or returns a HTTP 4xx or 5xx status in response to some or all of the dCDN's CDNI metadata requests, the dCDN MUST NOT serve the requested content.

Where a dCDN is interconnected with multiple uCDNs, the dCDN needs to determine which uCDN's CDNI metadata should be used to handle a particular User Agent request.

When application level redirection (e.g., HTTP 302 redirects) is being used between CDNs, it is expected that the dCDN will be able to determine the uCDN that redirected a particular request from information contained in the received request (e.g., via the URI). With knowledge of which uCDN routed the request, the dCDN can choose the correct uCDN from which to obtain the HostIndex. Note that the HostIndexes served by each uCDN can be unique.

In the case of DNS redirection there is not always sufficient information carried in the DNS request from User Agents to determine the uCDN that redirected a particular request (e.g., when content from a given host is redirected to a given dCDN by more than one uCDN) and therefore dCDNs will have to apply local policy when deciding which uCDN's metadata to apply.

6.3. Bootstrapping

The URI for the HostIndex object of a given uCDN needs to be configured in the dCDN. All other objects/resources are then discoverable from the HostIndex object by following any links in the HostIndex object and through the referenced HostMetadata and PathMetadata objects and their GenericMetadata sub-objects.

Manual configuration of the URI for the HostIndex object is outside the scope of this document.

6.4. Encoding

CDNI metadata objects MUST be encoded as I-JSON objects [RFC7493] containing a dictionary of (key,value) pairs where the keys are the property names and the values are the associated property values.

The keys of the dictionary are the names of the properties associated with the object and are therefore dependent on the specific object being encoded (i.e., dependent on the CDNI Payload Type of the returned resource). Likewise, the values associated with each

property (dictionary key) are dependent on the specific object being encoded (i.e., dependent on the CDNI Payload Type of the returned resource).

Dictionary keys (properties) in I-JSON are case sensitive. By convention, any dictionary key (property) defined by this document (for example, the names of CDNI metadata object properties) MUST be lowercase.

6.5. Extensibility

The set of GenericMetadata objects can be extended with additional (standards based or vendor specific) metadata objects through the specification of new GenericMetadata objects. The GenericMetadata object defined in Section 4.1.7 specifies a type field and a type-specific value field that allows any metadata to be included in either the HostMetadata or PathMetadata arrays.

As with the initial GenericMetadata types defined in Section 4.2, future GenericMetadata types MUST specify the information necessary for constructing and decoding the GenericMetadata object.

Any document which defines a new GenericMetadata type MUST:

1. Specify and register the CDNI Payload Type [RFC7736] used to identify the new GenericMetadata type being specified.
2. Define the set of properties associated with the new GenericMetadata object. GenericMetadata MUST NOT contain a property named "href" because doing so would conflict with the ability to detect Link objects (see Section 4.3.1).
3. Define a name, description, type, and whether or not the property is mandatory-to-specify.
4. Describe the semantics of the new type including its purpose and example of a use case to which it applies including an example encoded in I-JSON.
5. Describe the security and privacy consequences, for both the user-agent and the CDN, of the new GenericMetadata object.
6. Describe any relation to, conflict with, or obsolescence of other existing CDNI metadata objects.

Note: In the case of vendor specific extensions, vendor-identifying CDNI Payload Type names will decrease the possibility of GenericMetadata type collisions.

6.6. Metadata Enforcement

At any given time, the set of GenericMetadata types supported by the uCDN might not match the set of GenericMetadata types supported by the dCDN.

In cases where a uCDN sends metadata containing a GenericMetadata type that a dCDN does not support, the dCDN MUST enforce the semantics of the "mandatory-to-enforce" property. If a dCDN does not understand or is unable to perform the functions associated with any "mandatory-to-enforce" metadata, the dCDN MUST NOT service any requests for the corresponding content.

Note: Ideally, uCDNs would not delegate content requests to a dCDN that does not support the "mandatory-to-enforce" metadata associated with the content being requested. However, even if the uCDN has a priori knowledge of the metadata supported by the dCDN (e.g., via the FCI or through out-of-band negotiation between CDN operators), metadata support can fluctuate or be inconsistent (e.g., due to miscommunication, mis-configuration, or temporary outage). Thus, the dCDN MUST always evaluate all metadata associated with redirection and content requests and reject any requests where "mandatory-to-enforce" metadata associated with the content cannot be enforced.

6.7. Metadata Conflicts

It is possible that new metadata definitions will obsolete or conflict with existing GenericMetadata (e.g., a future revision of the CDNI metadata interface could redefine the Auth GenericMetadata object or a custom vendor extension could implement an alternate Auth metadata option). If multiple metadata (e.g., MI.Auth.v2, vendor1.Auth, and vendor2.Auth) all conflict with an existing GenericMetadata object (i.e., MI.Auth) and all are marked as "mandatory-to-enforce", it could be ambiguous which metadata should be applied, especially if the functionality of the metadata overlap.

As described in Section 3.3, metadata override only applies to metadata objects of the same exact type found in HostMetadata and nested PathMetadata structures. The CDNI metadata interface does not support enforcement of dependencies between different metadata types. It is the responsibility of the CSP and the CDN operators to ensure that metadata assigned to a given piece of content do not conflict.

Note: Because metadata is inherently ordered in HostMetadata and PathMetadata arrays, as well as in the PathMatch hierarchy, multiple conflicting metadata types MAY be used, however, metadata hierarchies SHOULD ensure that independent PathMatch root objects are used to prevent ambiguous or conflicting metadata definitions.

6.8. Versioning

The version of CDNI metadata objects is conveyed inside the CDNI Payload Type that is included in either the HTTP Content-Type header, for example: "Content-Type: application/cdni; ptype=MI.HostIndex", when retrieved via a link, or in the link type (Section 4.3.1), generic-metadata-type (Section 4.1.7), or auth-type (Section 4.2.7) properties in the JSON payload. The CDNI Payload Type uniquely identifies the specification defining that object, including any relation to, conflicts with, or obsolescence of other metadata. There is no explicit version mapping requirement, however, for ease of understanding, metadata creators SHOULD make new versions of metadata easily visible via the CDNI Payload Type, e.g., by appending a version string. Note: A version string is optional on the first version, e.g., MI.HostIndex, but could be added for subsequent versions, e.g., MI.HostIndex.v2, MI.HostIndex.v3, etc.

Except when referenced by a Link object, nested metadata objects (i.e., structural metadata below the HostIndex; Source objects; Location, TimeWindow, and Protocol Rule objects; and Footprint and TimeWindow objects) can be serialized into a JSON payload without explicit CDNI Payload Type information. The type is inferred from the outer structural metadata, generic metadata, or auth object CDNI Payload Type. To avoid ambiguity when revising nestable metadata objects, any outer metadata object(s) MUST be reversioned and allocated new CDNI Payload Type(s) at the same time. For example, the MI.HostIndex object defined in this document contains an array of MI.HostMatch objects, which each in turn contains a MI.HostMetadata object. If a new MI.HostMetadata.v2 object were required, the outer MI.HostIndex and MI.HostMatch objects would need to be revised, e.g., to MI.HostIndex.v2 and MI.HostMatch.v2, respectively. Similarly, if a new MI.TimeWindowRule.v2 object was required, the outer MI.TimeWindowACL object would need to be revised, e.g., to MI.TimeWindowACL.v2; the MI.TimeWindowRule.v2 object, though, could still contain MI.TimeWindow objects, if so specified.

HTTP requests sent to a metadata server SHOULD include an Accept header with the CDNI Payload Type of the expected object. Metadata clients can specify multiple CDNI Payload Types in the Accept header, for example, if a metadata client is capable of processing two different versions of the same type of object (defined by different CDNI Payload Types) it might decide to include both in the Accept header.

6.9. Media Types

All CDNI metadata objects use the Media Type "application/cdni". The CDNI Payload Type for each object then contains the object name of that object as defined by this document, prefixed with "MI.". Table 4 lists the CDNI Payload Type for the metadata objects (resources) specified in this document.

Data Object	CDNI Payload Type
HostIndex	MI.HostIndex
HostMatch	MI.HostMatch
HostMetadata	MI.HostMetadata
PathMatch	MI.PathMatch
PatternMatch	MI.PatternMatch
PathMetadata	MI.PathMetadata
SourceMetadata	MI.SourceMetadata
Source	MI.Source
LocationACL	MI.LocationACL
LocationRule	MI.LocationRule
Footprint	MI.Footprint
TimeWindowACL	MI.TimeWindowACL
TimeWindowRule	MI.TimeWindowRule
TimeWindow	MI.TimeWindow
ProtocolACL	MI.ProtocolACL
ProtocolRule	MI.ProtocolRule
DeliveryAuthorization	MI.DeliveryAuthorization
Cache	MI.Cache
Auth	MI.Auth
Grouping	MI.Grouping

Table 4: CDNI Payload Types for CDNI Metadata objects

6.10. Complete CDNI Metadata Example

A dCDN requests the HostIndex and receive the following object with a CDNI payload type of "MI.HostIndex":

```

{
  "hosts": [
    {
      "host": "video.example.com",
      "host-metadata": {
        "type": "MI.HostMetadata",
        "href": "https://metadata.ucdn.example/host1234"
      }
    },
    {
      "host": "images.example.com",
      "host-metadata": {
        "type": "MI.HostMetadata",
        "href": "https://metadata.ucdn.example/host5678"
      }
    }
  ]
}

```

If the incoming request has a Host header with "video.example.com" then the dCDN would fetch the HostMetadata object from "https://metadata.ucdn.example/host1234" expecting a CDNI payload type of "MI.HostMetadata":

```

{
  "metadata": [
    {
      "generic-metadata-type": "MI.SourceMetadata",
      "generic-metadata-value": {
        "sources": [
          {
            "endpoint": ["acq1.ucdn.example"],
            "protocol": "http/1.1"
          },
          {
            "endpoint": ["acq2.ucdn.example"],
            "protocol": "http/1.1"
          }
        ]
      }
    },
    {
      "generic-metadata-type": "MI.LocationACL",
      "generic-metadata-value": {
        "locations": [
          {
            "footprints": [
              {

```



```

        "footprint-type": "ipv4cidr",
        "footprint-value": ["192.0.2.0/24"]
    },
    {
        "footprint-type": "ipv6cidr",
        "footprint-value": ["2001:db8::/32"]
    },
    {
        "footprint-type": "countrycode",
        "footprint-value": ["us"]
    },
    {
        "footprint-type": "asn",
        "footprint-value": ["as64496"]
    }
],
"action": "deny"
}
]
}
},
{
    "generic-metadata-type": "MI.ProtocolACL",
    "generic-metadata-value": {
        "protocol-acl": [
            {
                "protocols": [
                    "http/1.1"
                ],
                "action": "allow"
            }
        ]
    }
}
],
"paths": [
    {
        "path-pattern": {
            "pattern": "/video/trailers/*"
        },
        "path-metadata": {
            "type": "MI.PathMetadata",
            "href": "https://metadata.ucdn.example/host1234/pathABC"
        }
    },
    {
        "path-pattern": {
            "pattern": "/video/movies/*"

```

```

    },
    "path-metadata": {
      "type": "MI.PathMetadata",
      "href": "https://metadata.ucdn.example/host1234/pathDEF"
    }
  ]
}

```

Suppose the path of the requested resource matches the `"/video/movies/*"` pattern, the next metadata requested would be for `"https://metadata.ucdn.example/host1234/pathDCE"` with an expected CDNI payload type of `"MI.PathMetadata"`:

```

{
  "metadata": [],
  "paths": [
    {
      "path-pattern": {
        "pattern": "/videos/movies/hd/*"
      },
      "path-metadata": {
        "type": "MI.PathMetadata",
        "href": "https://metadata.ucdn.example/host1234/pathDEF/path123"
      }
    }
  ]
}

```

Finally, if the path of the requested resource also matches the `"/videos/movies/hd/*"` pattern, the dCDN would also fetch the following object from `"https://metadata.ucdn.example/host1234/pathDEF/path123"` with CDNI payload type `"MI.PathMetadata"`:

```
{
  "metadata": [
    {
      "generic-metadata-type": "MI.TimeWindowACL",
      "generic-metadata-value": {
        "times": [
          "windows": [
            {
              "start": "1213948800",
              "end": "1327393200"
            }
          ],
          "action": "allow"
        ]
      }
    ]
  }
}
```

The final set of metadata which applies to the requested resource includes a SourceMetadata, a LocationACL, a ProtocolACL, and a TimeWindowACL.

7. IANA Considerations

7.1. CDNI Payload Types

This document requests the registration of the following CDNI Payload Types under the IANA CDNI Payload Type registry:

Payload Type	Specification
MI.HostIndex	RFCthis
MI.HostMatch	RFCthis
MI.HostMetadata	RFCthis
MI.PathMatch	RFCthis
MI.PatternMatch	RFCthis
MI.PathMetadata	RFCthis
MI.SourceMetadata	RFCthis
MI.Source	RFCthis
MI.LocationACL	RFCthis
MI.LocationRule	RFCthis
MI.Footprint	RFCthis
MI.TimeWindowACL	RFCthis
MI.TimeWindowRule	RFCthis
MI.TimeWindow	RFCthis
MI.ProtocolACL	RFCthis
MI.ProtocolRule	RFCthis
MI.DeliveryAuthorization	RFCthis
MI.Cache	RFCthis
MI.Auth	RFCthis
MI.Grouping	RFCthis

[RFC Editor: Please replace RFCthis with the published RFC number for this document.]

7.1.1. CDNI MI HostIndex Payload Type

Purpose: The purpose of this payload type is to distinguish HostIndex MI objects (and any associated capability advertisement)

Interface: MI/FCI

Encoding: see Section 4.1.1

7.1.2. CDNI MI HostMatch Payload Type

Purpose: The purpose of this payload type is to distinguish HostMatch MI objects (and any associated capability advertisement)

Interface: MI/FCI

Encoding: see Section 4.1.2

7.1.3. CDNI MI HostMetadata Payload Type

Purpose: The purpose of this payload type is to distinguish HostMetadata MI objects (and any associated capability advertisement)

Interface: MI/FCI

Encoding: see Section 4.1.3

7.1.4. CDNI MI PathMatch Payload Type

Purpose: The purpose of this payload type is to distinguish PathMatch MI objects (and any associated capability advertisement)

Interface: MI/FCI

Encoding: see Section 4.1.4

7.1.5. CDNI MI PatternMatch Payload Type

Purpose: The purpose of this payload type is to distinguish PatternMatch MI objects (and any associated capability advertisement)

Interface: MI/FCI

Encoding: see Section 4.1.5

7.1.6. CDNI MI PathMetadata Payload Type

Purpose: The purpose of this payload type is to distinguish PathMetadata MI objects (and any associated capability advertisement)

Interface: MI/FCI

Encoding: see Section 4.1.6

7.1.7. CDNI MI SourceMetadata Payload Type

Purpose: The purpose of this payload type is to distinguish SourceMetadata MI objects (and any associated capability advertisement)

Interface: MI/FCI

Encoding: see Section 4.2.1

7.1.1.8. CDNI MI Source Payload Type

Purpose: The purpose of this payload type is to distinguish Source MI objects (and any associated capability advertisement)

Interface: MI/FCI

Encoding: see Section 4.2.1.1

7.1.1.9. CDNI MI LocationACL Payload Type

Purpose: The purpose of this payload type is to distinguish LocationACL MI objects (and any associated capability advertisement)

Interface: MI/FCI

Encoding: see Section 4.2.2

7.1.1.10. CDNI MI LocationRule Payload Type

Purpose: The purpose of this payload type is to distinguish LocationRule MI objects (and any associated capability advertisement)

Interface: MI/FCI

Encoding: see Section 4.2.2.1

7.1.1.11. CDNI MI Footprint Payload Type

Purpose: The purpose of this payload type is to distinguish Footprint MI objects (and any associated capability advertisement)

Interface: MI/FCI

Encoding: see Section 4.2.2.2

7.1.1.12. CDNI MI TimeWindowACL Payload Type

Purpose: The purpose of this payload type is to distinguish TimeWindowACL MI objects (and any associated capability advertisement)

Interface: MI/FCI

Encoding: see Section 4.2.3

7.1.13. CDNI MI TimeWindowRule Payload Type

Purpose: The purpose of this payload type is to distinguish TimeWindowRule MI objects (and any associated capability advertisement)

Interface: MI/FCI

Encoding: see Section 4.2.3.1

7.1.14. CDNI MI TimeWindow Payload Type

Purpose: The purpose of this payload type is to distinguish TimeWindow MI objects (and any associated capability advertisement)

Interface: MI/FCI

Encoding: see Section 4.2.3.2

7.1.15. CDNI MI ProtocolACL Payload Type

Purpose: The purpose of this payload type is to distinguish ProtocolACL MI objects (and any associated capability advertisement)

Interface: MI/FCI

Encoding: see Section 4.2.4

7.1.16. CDNI MI ProtocolRule Payload Type

Purpose: The purpose of this payload type is to distinguish ProtocolRule MI objects (and any associated capability advertisement)

Interface: MI/FCI

Encoding: see Section 4.2.4.1

7.1.17. CDNI MI DeliveryAuthorization Payload Type

Purpose: The purpose of this payload type is to distinguish DeliveryAuthorization MI objects (and any associated capability advertisement)

Interface: MI/FCI

Encoding: see Section 4.2.5

7.1.18. CDNI MI Cache Payload Type

Purpose: The purpose of this payload type is to distinguish Cache MI objects (and any associated capability advertisement)

Interface: MI/FCI

Encoding: see Section 4.2.6

7.1.19. CDNI MI Auth Payload Type

Purpose: The purpose of this payload type is to distinguish Auth MI objects (and any associated capability advertisement)

Interface: MI/FCI

Encoding: see Section 4.2.7

7.1.20. CDNI MI Grouping Payload Type

Purpose: The purpose of this payload type is to distinguish Grouping MI objects (and any associated capability advertisement)

Interface: MI/FCI

Encoding: see Section 4.2.8

7.2. CDNI Metadata Footprint Types Registry

The IANA is requested to create a new "CDNI Metadata Footprint Types" subregistry in the "Content Delivery Networks Interconnection (CDNI) Parameters" registry. The "CDNI Metadata Footprint Types" namespace defines the valid Footprint object type values used by the Footprint object in Section 4.2.2.2. Additions to the Footprint type namespace conform to the "Specification Required" policy as defined in [RFC5226]. The designated expert will verify that new type definitions do not duplicate existing type definitions and prevent gratuitous additions to the namespace. New registrations are required to provide a clear description of how to interpret new footprint types.

The following table defines the initial Footprint Registry values:

Footprint Type	Description	Specification
ipv4cidr	IPv4 CIDR address block	RFCthis
ipv6cidr	IPv6 CIDR address block	RFCthis
asn	Autonomous System (AS) Number	RFCthis
countrycode	ISO 3166-1 alpha-2 code	RFCthis

[RFC Editor: Please replace RFCthis with the published RFC number for this document.]

7.3. CDNI Metadata Protocol Types Registry

The IANA is requested to create a new "CDNI Metadata Protocol Types" subregistry in the "Content Delivery Networks Interconnection (CDNI) Parameters" registry. The "CDNI Metadata Protocol Types" namespace defines the valid Protocol object values in Section 4.3.2, used by the SourceMetadata and ProtocolACL objects. Additions to the Protocol namespace conform to the "Specification Required" policy as defined in [RFC5226], where the specification defines the Protocol Type and the protocol to which it is associated. The designated expert will verify that new protocol definitions do not duplicate existing protocol definitions and prevent gratuitous additions to the namespace.

The following table defines the initial Protocol values corresponding to the HTTP and HTTPS protocols:

Protocol Type	Description	Type Specification	Protocol Specifications
http/1.1	Hypertext Transfer Protocol -- HTTP/1.1	RFCthis	RFC7230
https/1.1	HTTP/1.1 Over TLS	RFCthis	RFC7230, RFC2818

[RFC Editor: Please replace RFCthis with the published RFC number for this document.]

8. Security Considerations

8.1. Authentication and Integrity

A malicious metadata server, proxy server, or attacker, impersonating an authentic uCDN metadata interface without being detected, could provide false metadata to a dCDN that either:

- o Denies service for one or more pieces of content to one or more User Agents;
- o Directs dCDNs to contact malicious origin servers instead of the actual origin servers, and substitute legitimate content with malware or slanderous alternate content; or
- o Removes delivery restrictions (e.g., LocationACL, TimeWindowACL, ProtocolACL, or Auth metadata), allowing access to content that would otherwise be denied, and thus possibly violating license restrictions and incurring unwarranted delivery costs.

Unauthorized access to metadata could also enable a malicious metadata client to continuously issue metadata requests in order to overload a uCDN's metadata server(s).

Unauthorized access to metadata could further result in leakage of private information. A malicious metadata client could request metadata in order to gain access to origin servers, as well as information pertaining to content restrictions.

An implementation of the CDNI metadata interface MUST use mutual authentication and message authentication codes to prevent unauthorized access to and undetected modification of metadata (see Section 8.3).

8.2. Confidentiality and Privacy

Unauthorized viewing of metadata could result in leakage of private information. Content provider origin and policy information is conveyed through the CDNI metadata interface. A third party could intercept metadata transactions in order to gain access to origin servers, as well as information pertaining to content restrictions and usage patterns.

Note: The distribution of metadata by a uCDN to dCDNs could introduce privacy concerns for some content providers, e.g., dCDNs accepting content requests for a content provider's content might be able to obtain additional information and usage patterns relating to the users of a content provider's services. Content providers with concerns about divulging information to dCDNs can instruct their uCDN partners not to use CDNI when delivering their content.

An implementation of the CDNI metadata interface MUST use strong encryption to prevent unauthorized interception or monitoring of metadata (see Section 8.3).

8.3. Securing the CDNI Metadata interface

An implementation of the CDNI metadata interface MUST support TLS transport as per [RFC2818] and [RFC7230].

TLS MUST be used by the server-side (dCDN) and the client-side (uCDN) of the CDNI metadata interface, including authentication of the remote end, unless alternate methods are used for ensuring the security of the information in the CDNI metadata interface requests and responses (such as setting up an IPsec tunnel between the two CDNs or using a physically secured internal network between two CDNs that are owned by the same corporate entity).

The use of TLS for transport of the CDNI metadata interface messages allows:

- o The dCDN and uCDN to authenticate each other.

and, once they have mutually authenticated each other, it allows:

- o The dCDN and uCDN to authorize each other (to ensure they are transmitting/receiving CDNI metadata requests and responses from an authorized CDN);
- o CDNI metadata interface requests and responses to be transmitted with confidentiality; and
- o The integrity of the CDNI metadata interface requests and responses to be protected during the exchange.

When TLS is used, the general TLS usage guidance in [RFC7525] MUST be followed.

9. Acknowledgements

The authors would like to thank David Ferguson, Francois Le Faucheur, Jan Seedorf and Matt Miller for their valuable comments and input to this document.

10. Contributing Authors

[RFC Editor Note: Please move the contents of this section to the Authors' Addresses section prior to publication as an RFC.]

Grant Watson
Velocix (Alcatel-Lucent)
3 Ely Road
Milton, Cambridge CB24 6AA
UK

Email: gwatson@velocix.com

Kent Leung
Cisco Systems
3625 Cisco Way
San Jose, 95134
USA

Email: kleung@cisco.com

11. References

11.1. Normative References

- [ISO3166-1] The International Organization for Standardization, "Codes for the representation of names of countries and their subdivisions -- Part 1: Country codes", ISO 3166-1:2013, 2013.
- [POSIX] Institute of Electrical and Electronics Engineers, "Information Technology Portable Operating System Interface (POSIX) Part 1: System Application Program Interface (API) [C Language]", IEEE P1003.1, 1990.
- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, DOI 10.17487/RFC1034, November 1987, <<http://www.rfc-editor.org/info/rfc1034>>.
- [RFC1123] Braden, R., Ed., "Requirements for Internet Hosts - Application and Support", STD 3, RFC 1123, DOI 10.17487/RFC1123, October 1989, <<http://www.rfc-editor.org/info/rfc1123>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.

- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, DOI 10.17487/RFC3986, January 2005, <<http://www.rfc-editor.org/info/rfc3986>>.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<http://www.rfc-editor.org/info/rfc4291>>.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, DOI 10.17487/RFC5226, May 2008, <<http://www.rfc-editor.org/info/rfc5226>>.
- [RFC5890] Klensin, J., "Internationalized Domain Names for Applications (IDNA): Definitions and Document Framework", RFC 5890, DOI 10.17487/RFC5890, August 2010, <<http://www.rfc-editor.org/info/rfc5890>>.
- [RFC5891] Klensin, J., "Internationalized Domain Names in Applications (IDNA): Protocol", RFC 5891, DOI 10.17487/RFC5891, August 2010, <<http://www.rfc-editor.org/info/rfc5891>>.
- [RFC5952] Kawamura, S. and M. Kawashima, "A Recommendation for IPv6 Address Text Representation", RFC 5952, DOI 10.17487/RFC5952, August 2010, <<http://www.rfc-editor.org/info/rfc5952>>.
- [RFC6707] Niven-Jenkins, B., Le Faucheur, F., and N. Bitar, "Content Distribution Network Interconnection (CDNI) Problem Statement", RFC 6707, DOI 10.17487/RFC6707, September 2012, <<http://www.rfc-editor.org/info/rfc6707>>.
- [RFC7230] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", RFC 7230, DOI 10.17487/RFC7230, June 2014, <<http://www.rfc-editor.org/info/rfc7230>>.
- [RFC7493] Bray, T., Ed., "The I-JSON Message Format", RFC 7493, DOI 10.17487/RFC7493, March 2015, <<http://www.rfc-editor.org/info/rfc7493>>.
- [RFC7525] Sheffer, Y., Holz, R., and P. Saint-Andre, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", BCP 195, RFC 7525, DOI 10.17487/RFC7525, May 2015, <<http://www.rfc-editor.org/info/rfc7525>>.

11.2. Informative References

- [I-D.ietf-cdni-control-triggers]
Murray, R. and B. Niven-Jenkins, "CDNI Control Interface / Triggers", draft-ietf-cdni-control-triggers-15 (work in progress), May 2016.
- [I-D.ietf-cdni-redirection]
Niven-Jenkins, B. and R. Brandenburg, "Request Routing Redirection interface for CDN Interconnection", draft-ietf-cdni-redirection-20 (work in progress), August 2016.
- [I-D.ietf-cdni-uri-signing]
Leung, K., Faucheur, F., Brandenburg, R., Downey, B., and M. Fisher, "URI Signing for CDN Interconnection (CDNI)", draft-ietf-cdni-uri-signing-09 (work in progress), June 2016.
- [RFC2818] Rescorla, E., "HTTP Over TLS", RFC 2818, DOI 10.17487/RFC2818, May 2000, <<http://www.rfc-editor.org/info/rfc2818>>.
- [RFC6793] Vohra, Q. and E. Chen, "BGP Support for Four-Octet Autonomous System (AS) Number Space", RFC 6793, DOI 10.17487/RFC6793, December 2012, <<http://www.rfc-editor.org/info/rfc6793>>.
- [RFC7234] Fielding, R., Ed., Nottingham, M., Ed., and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Caching", RFC 7234, DOI 10.17487/RFC7234, June 2014, <<http://www.rfc-editor.org/info/rfc7234>>.
- [RFC7336] Peterson, L., Davie, B., and R. van Brandenburg, Ed., "Framework for Content Distribution Network Interconnection (CDNI)", RFC 7336, DOI 10.17487/RFC7336, August 2014, <<http://www.rfc-editor.org/info/rfc7336>>.
- [RFC7337] Leung, K., Ed. and Y. Lee, Ed., "Content Distribution Network Interconnection (CDNI) Requirements", RFC 7337, DOI 10.17487/RFC7337, August 2014, <<http://www.rfc-editor.org/info/rfc7337>>.
- [RFC7540] Belshe, M., Peon, R., and M. Thomson, Ed., "Hypertext Transfer Protocol Version 2 (HTTP/2)", RFC 7540, DOI 10.17487/RFC7540, May 2015, <<http://www.rfc-editor.org/info/rfc7540>>.

[RFC7736] Ma, K., "Content Delivery Network Interconnection (CDNI) Media Type Registration", RFC 7736, DOI 10.17487/RFC7736, December 2015, <<http://www.rfc-editor.org/info/rfc7736>>.

Authors' Addresses

Ben Niven-Jenkins
Velocix (Alcatel-Lucent)
3 Ely Road
Milton, Cambridge CB24 6AA
UK

Email: ben@velocix.com

Rob Murray
Velocix (Alcatel-Lucent)
3 Ely Road
Milton, Cambridge CB24 6AA
UK

Email: rmurray@velocix.com

Matt Caulfield
Cisco Systems
1414 Massachusetts Avenue
Boxborough, MA 01719
USA

Phone: +1 978 936 9307
Email: mcaulfie@cisco.com

Kevin J. Ma
Ericsson
43 Nagog Park
Acton, MA 01720
USA

Phone: +1 978-844-5100
Email: kevin.j.ma@ericsson.com

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 24, 2014

Danhua. Wang, Ed.
Huawei Technologies
B. Niven-Jenkins, Ed.
Velocix (Alcatel-Lucent)
Xiaoyan. He
Huawei
Chen. Ge
China Telecom
Wei. Ni
China Mobile
October 21, 2013

Request Routing Redirection Interface for CDN Interconnection
draft-ietf-cdni-redirection-01

Abstract

The Request Routing Interface comprises of (1) the asynchronous advertisement of footprint and capabilities by a downstream CDN that allows a upstream CDN to decide whether to redirect particular user requests to that downstream CDN; and (2) the synchronous operation of an upstream CDN requesting whether a downstream CDN is prepared to accept a user request and of a downstream CDN responding with how to actually redirect the user request. This document describes an interface for the latter part, i.e. the CDNI request routing/Redirection Interface.

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 April 24, 2014.

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	2
2. Terminology	3
3. Interface function and operation overview	4
4. HTTP based RESTful interface for the Redirection Interface .	5
4.1. Information passed in RI requests & responses	7
4.2. JSON encoding of RI requests & responses	9
4.3. MIME Media Types used by the RI interface	10
4.4. DNS redirection	10
4.4.1. DNS Redirection requests	10
4.4.2. DNS Redirection responses	12
4.5. HTTP Redirection	13
4.5.1. HTTP Redirection requests	13
4.5.2. HTTP Redirection responses	15
4.6. Indicating the cacheability and scope of responses . . .	16
4.7. Error responses	18
4.8. Loop detection & prevention	20
5. Security Considerations	20
6. IANA Considerations	21
7. Acknowledgements	21
8. Outstanding considerations	22
9. Contributing Authors	22
10. References	22
10.1. Normative References	22
10.2. Informative References	23
Authors' Addresses	23

1. Introduction

A Content Delivery Network (CDN) is a system built on an existing IP network which is used for large scale content delivery, via prefetching or dynamically caching content on its distributed surrogates (caching servers). [RFC6707] describes the problem area of interconnecting CDNs.

The CDNI request routing interface outlined in [I-D.ietf-cdni-framework] comprises of:

1. The asynchronous advertisement of footprint and capabilities by a downstream CDN that allows a upstream CDN to decide whether to redirect particular user requests to that downstream CDN.
2. The synchronous operation of an upstream CDN requesting whether a downstream CDN is prepared to accept a user request and of a downstream CDN responding with how to actually redirect the user request.

This document describes an interface for the latter part, i.e. the CDNI request routing/Redirection Interface (RI).

2. Terminology

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 [RFC2119].

This document reuses the terminology defined in [RFC6707]. The term "Distinguished CDN Domain" defined in [I-D.ietf-cdni-framework] is also reused in this document.

The following additional terms are introduced by this document:

Application Level Redirection: The act of using an application specific redirection mechanism for the request routing process of a CDN. The Redirection Target (RT) is the result of the routing decision of a CDN at the time it receives a content request via an application specific protocol response. Examples of an application level redirection are HTTP 302 Redirection and RTMP 302 Redirection.

DNS Redirection: The act of using DNS name resolution for the request routing process of a CDN. In DNS Redirection, the DNS name server of the CDN makes the routing decision based on a local policy and selects one or more Redirection Targets (RTs) and redirects the user agent to the RT(s) by returning the details of the RT(s) in response to the DNS query request from the user agent's DNS resolver.

HTTP Redirection: The act of using an HTTP redirection response for the request routing process of a CDN. The Redirection Target (RT) is the result of the routing decision of a CDN at the time it receives a content request via HTTP. HTTP Redirection is a particular case of Application Level Redirection.

Redirection Target (RT): A Redirection Target is the endpoint to which the user agent is redirected. In CDNI, a RT may point to a number of different components, some examples include a surrogate in the same CDN as the request router, a request router in a downstream CDN or a surrogate in a downstream CDN, etc.

3. Interface function and operation overview

[[Editor's note: Need to factor token authorisation into a future draft when that work is more stable/mature within the WG.]]

The CDNI request routing/Redirection Interface (RI) is one of the main building blocks required in order to interconnect CDNs. The main function of the Redirection Interface is to allow the Request Routing systems in interconnected CDNs to communicate to facilitate the redirection of User Agent requests between interconnected CDNs.

The detailed requirements for the Redirection Interface and their relative priorities are described in section 5 of [I-D.ietf-cdni-requirements].

The User Agent will make a request to a request router in the uCDN using one of either DNS or HTTP. If the RI is used between the uCDN and one or more dCDNs. The dCDN's RI response may contain a Redirection Target with a type that is compatible with the protocol used between User Agent and uCDN request router. The dCDN has control over the Redirection Target it provides and depending on the returned Redirection Target, the User Agent's request may be redirected to:

- o The final Surrogate, which may be in the dCDN or another dCDN (if dCDN delegates the delivery to another CDN).
- o A request router (in dCDN or another CDN) that will be using a redirection protocol (DNS or HTTP) which may or may not be the same as original redirection protocol.

The Redirection Interface operates between the Request Routing systems of a pair of interconnected CDNs. To enable communication over the Redirection Interface, the two interconnected CDNs need to know the end point (URI) in the other CDN to query. For example, an Upstream CDN needs to know the URI (end point) in a Downstream CDN to send its CDNI request routing queries to.

The Redirection Interface URI may be statically pre-configured, dynamically discovered via the CDNI control interface, or discovered via other means. However, such discovery mechanisms are not specified in this document, as they are considered out of the scope of the Redirection Interface specification.

CDNI solutions must support both of the request routing mechanisms illustrated in section 2.1 of [I-D.ietf-cdni-framework], namely Iterative Request Redirection and Recursive Request Redirection. However, the Iterative Request Redirection method does not invoke any interaction over the Redirection Interface between interconnected CDNs. Therefore, the Redirection Interface is only relevant in the case of Recursive Request Redirection and so this document will not discuss Iterative Request Redirection further.

In the case of Recursive Request Redirection, in order to perform redirection of a request received from a User Agent, the Upstream CDN queries the Downstream CDN so that the Downstream CDN can select and provide a Redirection Target. In cases where a uCDN has a choice of dCDNs it is down to the uCDN to decide (for example via configured policies) which dCDN(s) to query and in which order to query them. A number of strategies are possible including selecting a preferred dCDN based on local policy, possibly falling back to querying an alternative dCDN(s) if the first dCDN does not return a Redirection Target or otherwise reject the uCDN's RI request. A more complex strategy could be to query multiple dCDNs in parallel before selecting one and using the Redirection Target provided by that dCDN.

The Upstream CDN->User Agent redirection protocols addressed in this draft are: DNS redirection and HTTP redirection. Other types of application level redirection will not be discussed further in this draft. However the Redirection Interface is designed to be extensible and could be extended to support additional application level redirection protocols.

Also, according to the CDNI generic and request routing interface requirements, the CDNI solution shall support mechanisms to prevent and detect RI request loops. To meet such requirements, this document defines a loop prevention and detection mechanism as part of the Redirection Interface.

4. HTTP based RESTful interface for the Redirection Interface

This document defines a simple RESTful interface for the Redirection Interface based on HTTP [RFC2616], where the attributes of a User Agent's requests are encapsulated along with any other data that can aid the downstream CDN in processing the requests. The RI response encapsulates the attributes of the RT(s) that the upstream CDN should

return to the User Agent (if it decides to utilize the Downstream CDN for delivery) along with the policy for how the response can be reused.

The same RESTful interface is used for both DNS and HTTP redirection of User Agent's requests, although the contents of the RI requests/responses contain data specific to either DNS or HTTP redirection.

This approach has been chosen because it enables CDN operators to only have to deploy a single (RESTful) interface for the RI between their CDNs, regardless of the User Agent redirection method. In this way, from an operational point of view there is only one interface to monitor, manage, develop troubleshooting tools for, etc.

In addition, having a single RI where the attributes of the User Agent's DNS or HTTP request are encapsulated along with the other data required for the downstream CDN to make a request routing decision, avoids having to try and encapsulate or proxy DNS/HTTP/RTMP/etc requests and find ways to somehow embed the additional CDNI request routing/Redirection Interface properties/data within those End User DNS/HTTP/RTMP/etc requests.

Finally, the RI is easily extendable to support other User Agent request redirection methods (e.g. RTMP 302 redirection).

The generic Recursive Request Redirection message flow between Request Routing systems in a pair of interconnected CDNs is as follows:

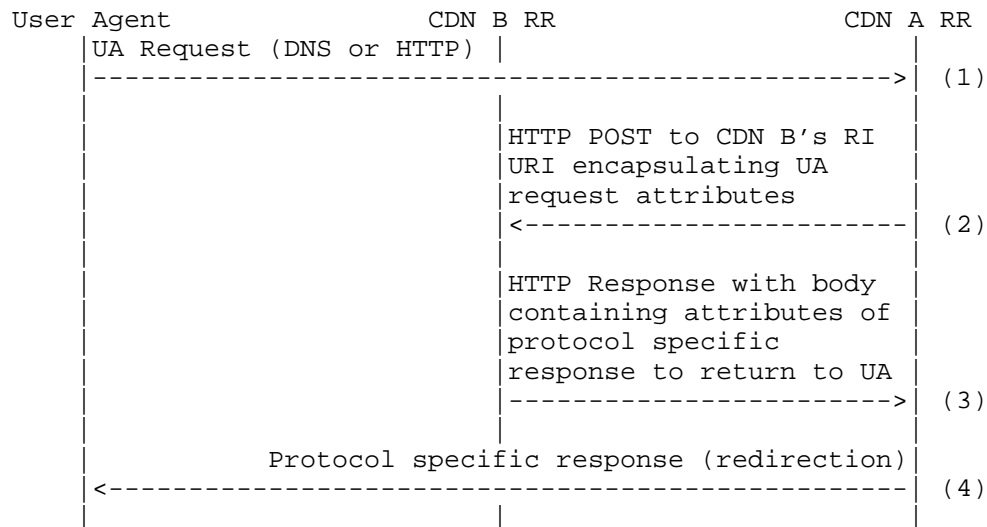


Figure 1: Generic Recursive Request Redirection message flow

1. The User Agent sends its request, either DNS request or HTTP request, to CDN A. The Request Routing System of CDN A processes the request and, through local policy, it recognizes that the request is best served by another CDN, specifically CDN B (or that CDN B is one of a number of candidate dCDNs it could use).
 2. The Request Routing System of CDN A sends an HTTP POST to CDN B's RI URI containing the attributes of the User Agent's request.
 3. The Request Routing System of CDN B processes the request and assuming the request is well formed, etc. responds with an HTTP "200" response with a message body containing the RT(s) to return to the User Agent as well as parameters that indicate the properties of the response (cacheability and scope).
 4. The Request Routing System of CDN A sends a protocol specific response (containing the returned attributes) to the User Agent, so that the User Agent's request will be redirected to the RT(s) returned by CDN B.
- 4.1. Information passed in RI requests & responses
- The information passed in RI requests splits into two basic categories:
1. The attributes of the User Agent's request to the upstream CDN.

2. Properties/parameters that the uCDN can use to control the dCDN's response or that can help the dCDN make its decision.

To assist the routing decision of a Downstream CDN, the Upstream CDN shall convey as much information as possible to the Downstream CDN, for example the URI of the requested content and the User Agent's location information, when those are known by the uCDN Request Routing system.

In order for the Downstream CDN to determine whether it is capable of delivering any requested content, it requires CDNI metadata related to the content the User Agent is requesting. That metadata will describe the content and any policies associated with it. It is expected that the RI request contains sufficient information for the Request Router in the Downstream CDN to be able to retrieve the require CDNI Metadata via the CDNI Metadata interface.

The information passed in RI responses splits into two basic categories:

1. The attributes of the RT to return to the User Agent in the DNS response or HTTP response.
2. Parameters/policies that indicate the properties of the response, such as, whether it is cacheable, the scope of the response, etc.

In addition to details of how to redirect the User Agent, the Downstream CDN may wish to return additional policy to the Upstream CDN to help the Upstream CDN with future RI requests. For example the Downstream CDN may wish to return a policy that expresses "this response can be reused without requiring a RI request for 60 seconds provided the User Agent's IP address is in the range 192.0.2.0 - 192.0.2.255".

These additional policies split into two basic categories:

- o An indication of the cacheability of the response carried in the HTTP response headers (to reduce the number of subsequent RI requests the uCDN needs to make).
- o The scope of the response (if it is cacheable) carried within the body of the HTTP response. For example whether the response applies to a wider range of IP addresses than what was included in the RI request.

The cacheability of the response is indicated using the standard HTTP Cache-Control mechanisms.

4.2. JSON encoding of RI requests & responses

The body of RI requests and responses is a JSON object containing a dictionary of keys. Keys MUST always be encoded in lowercase. Unknown keys MUST be ignored but the response MUST NOT be considered invalid unless the syntax of the request is invalid.

The following keys are defined:

Key	Request/Response	Description
dns	Both	The attributes of the UA's DNS request or the attributes of the RT(s) to return in a DNS response.
http	Both	The attributes of the UA's HTTP request or the attributes of the RT to return in a HTTP response.
scope	Response	The scope of the response (if it is cacheable). For example whether the response applies to a wider range of IP addresses than what was included in the RI request.
error	Response	Additional details if the response is an error response.
cdn-path	Both	A List of Strings. Contains the CDN Provider IDs of previous CDNs this RI request has passed through. When cascading a RI request the transit CDN appends its own CDN Provider ID to the list in cdn-path so that downstream CDNs can detect loops in the RI request chain. Transit CDNs should check the cdn-path and not cascade the RI request to downstream CDNs that are already listed in cdn-path. The cdn-path MUST be reflected back in RI responses.
max-hops	Request	Integer specifying the Maximum Number of hops (CDN Provider IDs) this request is allowed to be propagated along. This allows the uCDN to crudely constrain the latency of the request routing


```

|               |               | chain.               |
+-----+-----+-----+-----+

```

Top-Level keys in RI requests/responses

A single request or response MUST contain only one of the dns or http keys. Requests MUST contain a cdn-path key.

[[Editor's note: Need some text on minimum attributes to be able to (at least parse) - e.g. A/AAAA/CNAME, etc]]

Note: All implementations MUST support IPv4 addresses encoded as specified by the 'IPv4address' rule in Section 3.2.2 of [RFC3986] and MUST support all IPv6 address formats specified in [RFC4291]. Server implementations SHOULD use IPv6 address formats specified in [RFC5952].

4.3. MIME Media Types used by the RI interface

RI requests SHOULD use a MIME Media Type of application/cdni.redirectionrequest

RI responses SHOULD use a MIME Media Type of application/cdni.redirectionresponse.

4.4. DNS redirection

The following sections provide more detailed descriptions of the information that should be passed in RI requests and responses for DNS redirection.

4.4.1. DNS Redirection requests

For DNS based redirection the uCDN needs to pass the following information to the dCDN in the RI request:

- o The IP address of the DNS resolver that made the DNS request to the Upstream CDN.
- o The type of DNS query made (A, AAAA, RCODEs, etc.).
- o The class of DNS query made (usually IN). [[Editor's Note: Do we need to include class or can we always assume it is IN?]]
- o The fully qualified domain name for which DNS redirection is being requested.

- o The IP address or prefix of the User Agent (if known to the Upstream CDN, e.g. through draft-vandergaast-edns-client-subnet).

The information above is encoded as a set of key:value pairs within the dns dictionary as follows:

Key	Value	Mandatory	Description
resolver-ip	String	Yes	The IP address of the UA's DNS resolver.
qtype	String	Yes	The type of DNS query made by the UA's DNS resolvers in uppercase (A, AAAA, etc.).
qclass	String	Yes	The class of DNS query made in uppercase (IN, etc.).
qname	String	Yes	The fully qualified domain name being queried.
c-subnet	String	No	The IP address of the UA in CIDR format.
dns-only	Boolean	No	If True then dCDN MUST only use DNS redirection to a surrogate and MUST include the dns-only property set to True on any cascaded RI requests. Defaults to False.

The dns dictionary of a RI request for DNS based redirection MUST contain the following keys: resolver-ip, qtype, qclass, qname and the value of each MUST be the value of the appropriate part of the User Agent's DNS query/request.

An example RI request (uCDN->dCDN) for DNS based redirection:

```
POST /dcdn/ri HTTP/1.1
Host: rrl.dcdn.example.net
Accept: application/vnd.cdni.ri.response+json
```

```
{
  "dns" : {
    "resolver-ip" : "192.0.2.1",
    "c-subnet" : "198.51.100.0/24",
    "qtype" : "A",
    "qclass" : "IN",
```

```

    "qname" : "www.example.com"
  },
  "cdn-path": ["AS65551:0"],
  "max-hops": 3
}

```

4.4.2. DNS Redirection responses

For DNS based redirection the dCDN needs to return one of the following to the uCDN in the RI response:

- o The IP address of (or a CNAME to) the RT (if the dCDN is performing DNS based redirection); or
- o The IP address of (or a CNAME to) a RT which is a Request Router (if the dCDN is performing HTTP based redirection).

The information above is encoded as a set of key:value pairs within the dns dictionary as follows:

Key	Value	Mandatory	Description
rcode	Integer	Yes	DNS response code.
name	String	Yes	The fully qualified domain name the response relates to.
a	List of String	No	Set of IPv4 Addresses of RT(s).
aaaa	List of String	No	Set of IPv6 Addresses of RT(s).
cname	List of String	No	Set of fully qualified domain names of RT(s).
ttl	Integer	No	TTL of DNS response. Default is 0.

Response must contain at least one of a, aaaa, cname.

An example of a successful RI response (dCDN->uCDN) for DNS based redirection:

[[Editor's note: Currently shows both A/AAAA & CNAME in single response, need to split to show the different use cases]]

```

HTTP/1.1 200 OK
Date: Mon, 06 Aug 2012 18:41:38 GMT
Content-Type: application/vnd.cdni.rri.response+json

```

```

{
  "dns" : {
    "rcode" : 0,
    "name" : "www.example.com",
    "a" : ["192.0.2.200", "192.0.2.201"],
    "aaaa" : ["2001:DB8::C8", "2001:DB8::C9"],
    "cname" : ["rr1.dcdn.example",
               "rr2.dcdn.example"],
    "ttl" : 60
  }
}

```

4.5. HTTP Redirection

The following sections provide more detailed descriptions of the information that should be passed in RI requests and responses for HTTP redirection.

4.5.1. HTTP Redirection requests

For HTTP based redirection the uCDN MUST pass the following information to the dCDN in the RI request:

- o The IP address of the User Agent.
- o The URL requested by the User Agent.

The uCDN MAY also pass additional information to the dCDN in the RI request, such as:

- o The HTTP method or version number of the User Agent's request.
- o Additional HTTP header included in the User Agent request.

The information above is encoded as a set of key:value pairs within the http dictionary as follows:

Key	Value	Mandatory	Description
c-ip	String	Yes	The IP address of the UA.
cs-uri	String	Yes	The URI requested by

cs-method	String	Yes	the UA. The Method part of the Request-Line as defined in Section 5.1 of [RFC2616].
cs-version	String	Yes	The HTTP-Version part of the Request-Line as defined in Section 5.1 of [RFC2616].
cs(<HeaderName>)	String	No	The contents of the HTTP header named <HeaderName> as a string, for example cs(Cookie) would contain the content of the HTTP Cookie: header.

The http dictionary of a RI request for HTTP based redirection MUST contain the following keys: c-ip, cs-method, cs-version and cs-uri and the value of each MUST be the value of the appropriate part of the User Agent's DNS query/request.

An example RI request (uCDN->dCDN) for HTTP based redirection:

```
POST/dcdn/rrri HTTP/1.1
Host: rr1.dcdn.example.net
Accept: application/vnd.cdni.rrri.response+json
```

```
{
  "http": {
    "c-ip": "198.51.100.1",
    "cs-uri": "http://www.example.com",
    "cs-version": "HTTP/1.1",
    "cs-method": "GET"
  },
  "cdn-path": ["AS65551:0"],
  "max-hops": 3
}
```

4.5.2. HTTP Redirection responses

For HTTP based redirection the dCDN needs to return one of the following to the uCDN in the RI response:

- o A URL pointing to the selected RT (if the dCDN is redirecting the User Agent directly to a surrogate); or
- o A URL pointing to a RT which is a Request Router (if the dCDN is not redirecting the User Agent directly to a surrogate).

The information above is encoded as a set of key:value pairs within the http dictionary as follows:

Key	Value	Mandatory	Description
sc-status	Integer	Yes	The status code of the HTTP response to return to the UA (usually 302).
cs-uri	String	Yes	The URI requested by the UA/client.
sc-location	String	Yes	The contents of the Location header to return to the UA (i.e. a URI pointing to the RT(s)).
sc-cache-control	String	No	The contents of the Cache-Control header to return to the UA.

[[Editor's Note: Should we change the format above to align with the cs() format for headers on the RI request and allow the dCDN to signal back any headers it wants in the response as sc(<HeaderName>)? How to handle sc-status in that case - as a "special" header or separate key? Probably need to give some advice on HTTP headers the uCDN may want to override/not pass through, e.g. Server:?]]

An example of a successful RI response (dCDN->uCDN) for HTTP based redirection:

```
HTTP/1.1 200 OK
Date: Mon, 06 Aug 2012 18:41:38 GMT
Content-Type: application/vnd.cdni.ri.response+json
```

```

{
  "http": {
    "sc-status": 302,
    "cs-uri": "http://www.example.com"
    "sc-location":
      "http://surl.dcdn.example/ucdn/example.com",
    "sc-cache-control" : "public, max-age=30"
  }
}

```

4.6. Indicating the cacheability and scope of responses

RI responses may be cacheable and may be reused by the uCDN in response to User Agent requests without the uCDN issuing another RI request to the dCDN if the RI response is considered cacheable & not stale according to the standard HTTP Cache-Control, etc mechanisms. Additionally, an RI response MUST NOT be reused unless the request from the User Agent would generate an identical RI request to the dCDN as the one that resulted in the cached RI response.

Additionally, although RI requests only encode a single User Agent request to be redirected there may be cases where a dCDN wishes to indicate to the uCDN that the RI response can be reused for other User Agent requests without the uCDN having to make another request via the RI. For example a dCDN may know that it will always select the same Surrogates for a given set of User Agent IP addresses and in order to reduce request volume across the RI or to remove the additional latency associated with an RI request, the dCDN may wish to indicate that set of User Agent IP addresses to the uCDN in the initial RI response. This is achieved by including an optional scope dictionary in the RI response.

Scope is encoded as a set of key:value pairs within the scope dictionary as follows:

Key	Value	Mandatory	Description
iprange	List of String	No	A List of IP subnets in CIDR notation that this RI response can be reused for, provided the RI response is still considered fresh.

If a uCDN has multiple cached responses with overlapping scopes, longest prefix matching of the User Agent's IP against the IP subnets in the scope of each response SHOULD be used to select the most appropriate RI response to use. [[Editor's note: is this always true? What about the most recent response, should that override older ones for the overlappign scope?]]

Example of DNS redirection response from Section 4.4.2 that is cacheable by the uCDN for 60 seconds and can be returned to any User Agent with an IPv4 address in 198.51.100.0/16.

```
HTTP/1.1 200 OK
Date: Mon, 06 Aug 2012 18:41:38 GMT
Content-Type: application/vnd.cdni.rri.response+json
Cache-Control: public, max-age=60
```

```
{
  "dns" : {
    "rcode" : 0,
    "name" : "www.example.com",
    "a" : ["192.0.2.200", "192.0.2.201"],
    "aaaa" : ["2001:DB8::C8", "2001:DB8::C9"],
    "cname" : ["rr1.dcdn.example",
               "rr2.dcdn.example"],
    "ttl" : 60
  }
  "scope" : {
    "iprange" : ["198.51.100.0/16"]
  }
}
```

Example of HTTP redirection response from Section 4.5.2 that is cacheable by the uCDN for 60 seconds and can be returned to any User Agent with an IPv4 address in 198.51.100.0/16.

```
HTTP/1.1 200 OK
Date: Mon, 06 Aug 2012 18:41:38 GMT
Content-Type: application/vnd.cdni.rri.response+json
Cache-Control: public, max-age=60
```

```
{
  "http": {
    "sc-status": 302,
    "cs-uri": "http://www.example.com"
    "sc-location":
      "http://surl.dcdn.example/ucdn/example.com",
    "sc-cache-control" : "public, max-age=30"
```



```

    }
    "scope" : {
      "iprange" : ["198.51.100.0/16"]
    }
  }
}

```

4.7. Error responses

[[Editor's note: Probably examples of errors that shouldn't be propagated to the User Agent?]]

Errors are indicated via an appropriate HTTP response code.

There will be many reasons that a server is not able to return a successful RI response to the client. To aid with diagnosing the cause of errors, RI responses may include an optional error dictionary to provide additional information to the client as to the reason/cause of the error. The intention behind the error dictionary is to aid with manual diagnostics of issues and not to provide an exhaustive, pre-specified, set of error conditions and associated RI specific error codes, etc.

[[Editor's note: We've tried to keep error specification light weight & provide the hooks needed to help with debugging without trying to be overly prescriptive over how it gets used as we'd like to avoid the rat hole of specifying every possible error condition and consequent actions. This is because it is hard to think of scenarios where having a set of pre-defined error codes/etc helps much because the relevant consequent action cannot be performed automatically by the client.]]

Additional error information (if present) is encoded as a set of key:value pairs within the error dictionary as follows:

Key	Value	Mandatory	Description
code	Integer	No	A numeric code defined by the server to indicate the error(s) that occurred.
description	String	No	A string providing further (probably human readable) information related to the error.

RI error response examples.

RI error response (dCDN->uCDN) for DNS based User Agent requests:

HTTP/1.1 500 Server Error

Date: Mon, 06 Aug 2012 18:41:38 GMT

Content-Type: application/vnd.cdni.rri.error+json

Cache-Control: private, no-cache

```
{
  "dns" : {
    "rcode" : 4                                # DNS response code (e.g.
                                              # doesn't support AAAA)
    "name" : "www.example.com",              # domain name response
                                              # relates to
  },
  "error" : {
    "code" : TBD,                             # Give each error type its
                                              # own numeric code
    "description" :                          # Give more informative
    "IPv6/AAAA queries are not supported"    # description than just
  }                                           # protocol specific error
                                              # codes
}
```

RI error response (dCDN->uCDN) for HTTP based User Agent requests:

HTTP/1.1 500 Server Error

Date: Mon, 06 Aug 2012 18:41:38 GMT

Content-Type: application/vnd.cdni.rri.error+json

Cache-Control: private, no-cache

```
{
  "http": {
    "sc-status": 400,                          # HTTP response code
    "url": "http://www.example.com",          # URL response
                                              # relates to
  },
  "error" : {
    "code" : TBD,                             # Give each error type its
                                              # own numeric code
    "description" : TBD                      # Give more informative
                                              # description than just
  }                                           # protocol specific error
                                              # codes
}
```

4.8. Loop detection & prevention

In order to prevent and detect RI request loops, each CDN MUST insert its CDN Provider ID into the cdn-path key of every RI request it originates or cascades. When receiving RI requests a dCDN should check the cdn-path and reject any RI requests which already contain the downstream CDN's Provider ID in the cdn-path. Transit CDNs should check the cdn-path and not cascade the RI request to downstream CDNs that are already listed in cdn-path. CDNs MUST NOT propagate to any downstream CDNs if the number of CDN Provider IDs in cdn-path (including the CDN's own Provider ID) is equal to or greater than max-hops.

The CDN Provider ID uniquely identifies each CDN provider during the course of request routing redirection. It consists of the characters AS followed by the CDN Provider's AS number, then a colon (':') and an additional qualifier that is used to guarantee uniqueness in case a particular AS has multiple independent CDNs deployed. For example "AS65551:0".

If a downstream CDN receives a RI request whose cdn-path already contains that downstream CDN's Provider ID the downstream CDN MUST send a RI response with an error code of [[TBD]].

It should be noted that the loop detection & prevention mechanisms described above only cover preventing and detecting loops within the RI itself. As well as loops with the RI itself, there is also the possibility of loops in the data plane, for example if the IP address(es) or URI(s) returned in RI responses do not resolve directly to a surrogate in the final dCDN there is the possibility that a User Agent may be continuously redirected through a loop of CDNs. The specification of solutions to address data plane request redirection loops between CDNs is out of the scope of this document.

5. Security Considerations

Information passed over the RI could be considered personal or sensitive, for example RI requests contain parts of a User Agent's original request and RI responses reveal information about the dCDN's policy for which surrogates should serve which content/user locations.

The RI interface also provides a mechanism whereby a uCDN could probe a dCDN and infer the dCDN's edge topology by making repeated RI requests for different content and/or UA IP addresses and correlating the responses from the dCDN. Additionally the ability for a dCDN to indicate that a RI response applies more widely than the original request (via the scope dictionary) may significantly reduce the

number of RI requests required to probe and infer the dCDN's edge topology.

The same information could be obtained in the absence of the RI interface, but it could be more difficult to gather as it would require a distributed set of machines with a range of different IP addresses each making requests directly to the dCDN. However, the RI facilitates easier collection of such information as it enables a single client to query the dCDN for a redirection/surrogate selection on behalf of any UA IP address.

In order to prevent passive interception of RI messages the RI communications channel should be suitably secured (e.g. use of TLS).

In order to reduce the risk of information leakage to unauthorized parties, RI clients and servers SHOULD use suitable authentication prior to trusting the contents of RI messages.

[[Editor's note: Not sure if the text below is really security considerations or whether it is better placed elsewhere in the document.]]

In HTTP based Recursive Request Redirection, the end user's web browsers will not send cookies if the content request is redirected to a URL in a different domain rather than the original CP's domain, e.g. the Downstream CDN's domain. If the browser is expected to send any cookies associated with the original CP's domain, this will cause problem that the CP's policy is not enforced by the CDN.

The section 5.2 of draft [I-D.peterson-cdni-strawman] has discussed a similar question and given a solution.

6. IANA Considerations

[[Editors' Note: Need to insert some text to register the Media Types we use with IANA?]]

7. Acknowledgements

The authors would like to thank Ray Brandenburg, Taesang Choi, Francois le Faucheur and Scott Wainner for their valuable comments and input to this document.

8. Outstanding considerations

Along with the various Editor's notes in the document, the following items still need to be addressed:

- o Additional examples of RI requests/responses (including error responses) illustrating different use cases.
- o Description/specification for how to extend the protocol with additional optional parameters/attributes.

9. Contributing Authors

[RFC Editor Note: Please move the contents of this section to the Authors' Addresses section prior to publication as an RFC.]

Spencer Dawkins
Huawei

Email: spencer@wonderhamster.org

Yunfei Zhang

Email: hishigh@gmail.com

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.
- [RFC5952] Kawamura, S. and M. Kawashima, "A Recommendation for IPv6 Address Text Representation", RFC 5952, August 2010.
- [RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2616, June 1999.

10.2. Informative References

- [RFC6707] Niven-Jenkins, B., Le Faucheur, F., and N. Bitar, "Content Distribution Network Interconnection (CDNI) Problem Statement", RFC 6707, September 2012.
- [I-D.ietf-cdni-framework]
Peterson, L. and B. Davie, "Framework for CDN Interconnection", draft-ietf-cdni-framework-03 (work in progress), February 2013.
- [I-D.ietf-cdni-requirements]
Leung, K. and Y. Lee, "Content Distribution Network Interconnection (CDNI) Requirements", draft-ietf-cdni-requirements-06 (work in progress), April 2013.
- [I-D.peterson-cdni-strawman]
Peterson, L. and J. Hartman, "A Simple Approach to CDN Interconnection", draft-peterson-cdni-strawman-01 (work in progress), May 2011.

Authors' Addresses

Wang Danhua (editor)
Huawei Technologies
No. 101 Software Avenue
Nanjing, Jiangsu Province 210001
P.R.China

Phone: +86-25-56624734
Fax: +86-25-56624702
Email: wangdanhua@huawei.com

Ben Niven-Jenkins (editor)
Velocix (Alcatel-Lucent)
3 Ely Road
Milton, Cambridge CB24 6DD
UK

Email: ben@velocix.com

He Xiaoyan
Huawei
B2, Huawei Industrial Base
518129
P.R.China

Email: hexiaoyan@huawei.com

Ge Chen
China Telecom
109 West Zhongshan Ave, Tianhe District
Guangzhou
P.R. China

Email: cheng@gsta.com

Ni Wei
China Mobile
No.32 Xuanwumen West Street Xicheng District
Beijing 100053
P.R. China

Email: niwei@chinamobile.com

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: February 16, 2017

B. Niven-Jenkins, Ed.
Nokia
R. van Brandenburg, Ed.
TNO
August 15, 2016

Request Routing Redirection interface for CDN Interconnection
draft-ietf-cdni-redirection-20

Abstract

The Request Routing Interface comprises (1) the asynchronous advertisement of footprint and capabilities by a downstream Content Delivery Network (CDN) that allows an upstream CDN to decide whether to redirect particular user requests to that downstream CDN; and (2) the synchronous operation of an upstream CDN requesting whether a downstream CDN is prepared to accept a user request and of a downstream CDN responding with how to actually redirect the user request. This document describes an interface for the latter part, i.e., the CDNI Request Routing Redirection interface.

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 February 16, 2017.

Copyright Notice

Copyright (c) 2016 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	2
2. Terminology	3
3. Interface function and operation overview	4
4. HTTP based interface for the Redirection Interface	5
4.1. Information passed in RI requests & responses	7
4.2. JSON encoding of RI requests & responses	8
4.3. MIME Media Types used by the RI interface	10
4.4. DNS redirection	10
4.4.1. DNS Redirection requests	10
4.4.2. DNS Redirection responses	12
4.5. HTTP Redirection	14
4.5.1. HTTP Redirection requests	14
4.5.2. HTTP Redirection responses	16
4.6. Cacheability and scope of responses	18
4.7. Error responses	20
4.8. Loop detection & prevention	24
5. Security Considerations	25
5.1. Authentication, Authorization, Confidentiality, Integrity Protection	26
5.2. Privacy	26
6. IANA Considerations	27
6.1. CDNI Payload Type Parameter registrations	27
6.1.1. CDNI RI Redirection Request Payload Type	27
6.1.2. CDNI RI Redirection Response Payload Type	28
6.2. RI Error response registry	28
7. Contributors	29
8. Acknowledgements	29
9. References	29
9.1. Normative References	29
9.2. Informative References	31
Authors' Addresses	31

1. Introduction

A Content Delivery Network (CDN) is a system built on an existing IP network which is used for large scale content delivery, via prefetching or dynamically caching content on its distributed surrogates (caching servers). [RFC6707] describes the problem area of interconnecting CDNs.

The CDNI Request Routing interface outlined in [RFC7336] comprises of:

1. The asynchronous advertisement of footprint and capabilities by a downstream CDN (dCDN) that allows an upstream CDN (uCDN) to decide whether to redirect particular user requests to that dCDN.
2. The synchronous operation of a uCDN requesting whether a dCDN is prepared to accept a user request and of a dCDN responding with how to actually redirect the user request.

This document describes an interface for the latter part, i.e., the CDNI Request Routing Redirection interface (RI).

2. Terminology

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 [RFC2119].

This document reuses the terminology defined in [RFC6707].

The following additional terms are introduced by this document:

Application Level Redirection: The act of using an application specific redirection mechanism for the request routing process of a CDN. The Redirection Target (RT) is the result of the routing decision of a CDN at the time it receives a content request via an application specific protocol response. Examples of an application level redirection are HTTP 302 Redirection and RTMP 302 Redirection [RTMP].

DNS Redirection: The act of using DNS name resolution for the request routing process of a CDN. In DNS Redirection, the DNS name server of the CDN makes the routing decision based on a local policy and selects one or more Redirection Targets (RTs) and redirects the user agent to the RT(s) by returning the details of the RT(s) in response to the DNS query request from the user agent's DNS resolver.

HTTP Redirection: The act of using an HTTP redirection response for the request routing process of a CDN. The Redirection Target (RT) is the result of the routing decision of a CDN at the time it receives a content request via HTTP. HTTP Redirection is a particular case of Application Level Redirection.

Redirection Target (RT): A Redirection Target is the endpoint to which the user agent is redirected. In CDNI, a RT may point to a

number of different components, some examples include a surrogate in the same CDN as the request router, a request router in a dCDN or a surrogate in a dCDN, etc.

3. Interface function and operation overview

The main function of the CDNI Redirection interface (RI) is to allow the request routing systems in interconnected CDNs to communicate to facilitate the redirection of User Agent requests between interconnected CDNs.

The detailed requirements for the Redirection interface and their relative priorities are described in section 5 of [RFC7337].

The User Agent will make a request to a request router in the uCDN using one of either DNS or HTTP. The RI is used between the uCDN and one or more dCDNs. The dCDN's RI response may contain a Redirection Target with a type that is compatible with the protocol used between User Agent and uCDN request router. The dCDN has control over the Redirection Target it provides. Depending on the returned Redirection Target, the User Agent's request may be redirected to:

- o The final Surrogate, which may be in the dCDN that returned the RI response to the uCDN, or another CDN (if the dCDN delegates the delivery to another CDN); or
- o A request router (in the dCDN or another CDN), which may use a different redirection protocol (DNS or HTTP) than the one included in the RI request.

The Redirection interface operates between the request routing systems of a pair of interconnected CDNs. To enable communication over the Redirection interface, the uCDN needs to know the URI (end point) in the dCDN to send CDNI request routing queries.

The Redirection interface URI may be statically pre-configured, dynamically discovered via the CDNI Control interface, or discovered via other means. However, such discovery mechanisms are not specified in this document, as they are considered out of the scope of the Redirection interface specification.

The Redirection interface is only relevant in the case of Recursive Request Redirection, as Iterative Request Redirection does not invoke any interaction over the Redirection interface between interconnected CDNs. Therefore, the scope of this document is limited to Recursive Request Redirection.

In the case of Recursive Request Redirection, in order to perform redirection of a request received from a User Agent, the uCDN queries the dCDN so that the dCDN can select and provide a Redirection Target. In cases where a uCDN has a choice of dCDNs it is up to the uCDN to decide (for example, via configured policies) which dCDN(s) to query and in which order to query them. A number of strategies are possible including selecting a preferred dCDN based on local policy, possibly falling back to querying an alternative dCDN(s) if the first dCDN does not return a Redirection Target or otherwise rejects the uCDN's RI request. A more complex strategy could be to query multiple dCDNs in parallel before selecting one and using the Redirection Target provided by that dCDN.

The uCDN->User Agent redirection protocols addressed in this draft are: DNS redirection and HTTP redirection. Other types of application level redirection will not be discussed further in this document. However, the Redirection interface is designed to be extensible and could be extended to support additional application level redirection protocols.

For both DNS & HTTP redirection, either HTTP or HTTPS could be used to connect to the Redirection Target. When HTTPS is used to connect to the uCDN, if the uCDN uses DNS redirection to identify the RT to the User Agent, then the new target domain name may not match the domain in the URL dereferenced to reach the uCDN; without operational precautions, and in the absence of DNSSEC, this can make a legitimate redirection look like a DNS-based attack to a User Agent and trigger security warnings. When DNS-based redirection with HTTPS is used, this specification assumes that any RT can complete the necessary TLS handshake with the User Agent. Any operational mechanisms this requires, e.g., private key distribution to surrogates and request routers in dCDNs, are outside the scope of this document.

This document also defines an RI loop prevention and detection mechanism as part of the Redirection interface.

4. HTTP based interface for the Redirection Interface

This document defines a simple interface for the Redirection interface based on HTTP [RFC7230], where the attributes of a User Agent's requests are encapsulated along with any other data that can aid the dCDN in processing the requests. The RI response encapsulates the attributes of the RT(s) that the uCDN should return to the User Agent (if it decides to utilize the dCDN for delivery) along with the policy for how the response can be reused. The examples of RI requests and responses below do not contain a complete set of HTTP headers for brevity; only the pertinent HTTP headers are shown.

The RI between the uCDN and dCDN uses the same HTTP interface to encapsulate the attributes of both DNS and HTTP requests received from User Agents, although the contents of the RI requests/responses contain data specific to either DNS or HTTP redirection.

This approach has been chosen because it enables CDN operators to only have to deploy a single interface for the RI between their CDNs, regardless of the User Agent redirection method. In this way, from an operational point of view there is only one interface to monitor, manage, develop troubleshooting tools for, etc.

In addition, having a single RI where the attributes of the User Agent's DNS or HTTP request are encapsulated along with the other data required for the dCDN to make a request routing decision, avoids having to try to encapsulate or proxy DNS/HTTP/RTMP/etc requests and find ways to somehow embed the additional CDNI Request Routing Redirection interface properties/data within those End User DNS/HTTP/RTMP/etc requests.

Finally, the RI is easily extendable to support other User Agent request redirection methods (e.g., RTMP 302 redirection) by defining additional protocol specific keys for RI requests and responses along with a specification how to process them.

The generic Recursive Request Redirection message flow between Request Routing systems in a pair of interconnected CDNs is as follows:

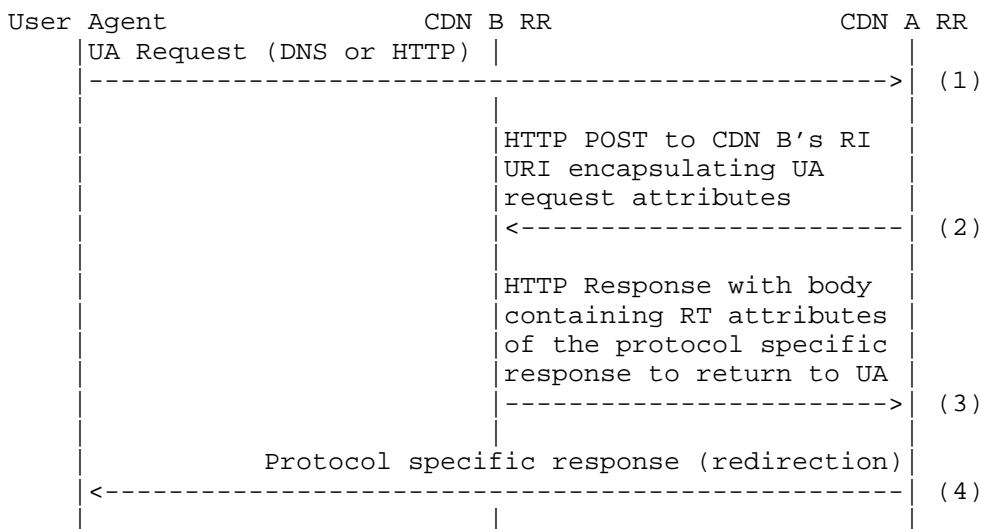


Figure 1: Generic Recursive Request Redirection message flow

1. The User Agent sends its (DNS or HTTP) request to CDN A. The Request Routing System of CDN A processes the request and, through local policy, recognizes that the request is best served by another CDN, specifically CDN B (or that CDN B may be one of a number of candidate dCDNs it could use).
2. The Request Routing System of CDN A sends an HTTP POST to CDN B's RI URI containing the attributes of the User Agent's request.
3. The Request Routing System of CDN B processes the RI request and assuming the request is well formed, responds with an HTTP "200" response with a message body containing the RT(s) to return to the User Agent as well as parameters that indicate the properties of the response (cacheability and scope).
4. The Request Routing System of CDN A sends a protocol specific response (containing the returned attributes) to the User Agent, so that the User Agent's request will be redirected to the RT(s) returned by CDN B.

4.1. Information passed in RI requests & responses

The information passed in RI requests splits into two basic categories:

1. The attributes of the User Agent's request to the uCDN.
2. Properties/parameters that the uCDN can use to control the dCDN's response or that can help the dCDN make its decision.

Generally, dCDNs can provide better routing decisions given additional information about the content request, e.g., the URI of the requested content or the User Agent's IP address or subnet. The set of information required to base a routing decision on can be highly dependent on the type of content delivered. A uCDN SHOULD only include information that is absolutely necessary for delivering that type of content. Cookies in particular are particularly sensitive from a security/privacy point of view and in general SHOULD NOT be conveyed in the RI Requests to the dCDN. The set of information necessary to be conveyed for a particular type of request is expected to be conveyed out of band between the uCDN and dCDN. See Section 5.2 for more detail on the privacy aspects of using RI Requests to convey information about UA requests.

In order for the dCDN to determine whether it is capable of delivering any requested content, it requires CDNI metadata related to the content the User Agent is requesting. That metadata will describe the content and any policies associated with it. It is

expected that the RI request contains sufficient information for the Request Router in the dCDN to be able to retrieve the required CDNI Metadata via the CDNI Metadata interface.

The information passed in RI responses splits into two basic categories:

1. The attributes of the RT to return to the User Agent in the DNS response or HTTP response.
2. Parameters/policies that indicate the properties of the response, such as, whether it is cacheable, the scope of the response, etc.

In addition to details of how to redirect the User Agent, the dCDN may wish to return additional policy information to the uCDN to it with future RI requests. For example, the dCDN may wish to return a policy that expresses "this response can be reused without requiring an RI request for 60 seconds provided the User Agent's IP address is in the range 198.51.100.0 - 198.51.100.255".

These additional policies split into two basic categories:

- o Cacheability information signaled via the HTTP response headers of the RI response (to reduce the number of subsequent RI requests the uCDN needs to make).
- o The scope of a cacheable response signaled in the HTTP response body of the RI response, for example, whether the response applies to a wider range of IP addresses than what was included in the RI request.

The cacheability of the response is indicated using the standard HTTP Cache-Control mechanisms.

4.2. JSON encoding of RI requests & responses

The body of RI requests and responses is a JSON object [RFC7159] that MUST conform to [RFC7493] containing a dictionary of key:value pairs. Senders MUST encode all (top level object and sub-object) keys specified in this document in lowercase. Receivers MUST ignore any keys that are unknown or invalid.

The following top level keys are defined along with whether they are applicable to RI requests, RI responses or both:

Key	Request/Response	Description
dns	Both	The attributes of the UA's DNS request or the attributes of the RT(s) to return in a DNS response.
http	Both	The attributes of the UA's HTTP request or the attributes of the RT to return in a HTTP response.
scope	Response	The scope of the response (if it is cacheable). For example, whether the response applies to a wider range of IP addresses than what was included in the RI request.
error	Response	Additional details if the response is an error response.
cdn-path	Both	A List of Strings. Contains a list of the CDN Provider IDs of previous CDNs that have participated in the request routing for the associated User Agent request. On RI requests it contains the list of previous CDNs that this RI request has passed through. On RI responses it contains the list of CDNs that were involved in obtaining the final redirection included in the RI response. See Section 4.8
max-hops	Request	Integer specifying the maximum number of hops (CDN Provider IDs) this request is allowed to be propagated along. This allows the uCDN to coarsely constrain the latency of the request routing chain.

Top-Level keys in RI requests/responses

A single request or response MUST contain only one of the dns or http keys. Requests MUST contain a cdn-path key and responses MAY contain a cdn-path key. If the max-hops key is not present then there is no limit on the number of CDN hops that the RI request can be propagated along. If the first uCDN does not wish the RI request to be propagated beyond the dCDN it is making the request to, then the uCDN MUST set max-hops to 1.

The `cdn-path` MAY be reflected back in RI responses, although doing so could expose information to the uCDN that a dCDN may not wish to expose (for example, the existence of business relationships between a dCDN and other CDNs).

If the `cdn-path` is reflected back in the RI response it MUST contain the value of `cdn-path` received in the associated RI request with the final dCDN's CDN Provider ID appended. Transit CDNs MAY remove the `cdn-path` from RI responses but MUST NOT modify the `cdn-path` in other ways.

The presence of an error key within a response that also contains either a `dns` or `http` key does not automatically indicate that the RI request was unsuccessful as the error key MAY be used for communicating additional (e.g., debugging) information. When a response contains an error key as well as either a `dns` or `http` key, the error-code SHOULD be `lxx` (e.g., 100). See Section 4.7 for more details of encoding error information in RI responses.

All implementations that support IPv4 addresses MUST support the encoding specified by the '`IPv4address`' rule in Section 3.2.2 of [RFC3986]. Likewise, implementations that support IPv6 addresses MUST support all IPv6 address formats specified in [RFC4291]. Server implementations SHOULD use IPv6 address formats specified in [RFC5952].

4.3. MIME Media Types used by the RI interface

RI requests MUST use a MIME Media Type of `application/cdni` as specified in [RFC7736], with the Payload Type (`ptype`) parameter set to '`redirection-request`'.

RI responses MUST use a MIME Media Type of `application/cdni` as specified in [RFC7736], with the Payload Type (`ptype`) parameter set to '`redirection-response`'.

4.4. DNS redirection

The following sections provide detailed descriptions of the information that should be passed in RI requests and responses for DNS redirection.

4.4.1. DNS Redirection requests

For DNS based redirection the uCDN needs to pass the following information to the dCDN in the RI request:

- o The IP address of the DNS resolver that made the DNS request to the uCDN.
- o The type of DNS query made (usually either A or AAAA).
- o The class of DNS query made (usually IN).
- o The fully qualified domain name for which DNS redirection is being requested.
- o The IP address or prefix of the User Agent (if known to the uCDN).

The information above is encoded as a set of key:value pairs within the dns dictionary as follows:

Key	Value	Mandatory	Description
resolver-ip	String	Yes	The IP address of the UA's DNS resolver.
qtype	String	Yes	The type of DNS query made by the UA's DNS resolvers in uppercase. The value of this field SHALL be set to either 'A' or 'AAAA'.
qclass	String	Yes	The class of DNS query made in uppercase (IN, etc.).
qname	String	Yes	The fully qualified domain name being queried.
c-subnet	String	No	The IP address (or prefix) of the UA in CIDR format.
dns-only	Boolean	No	If True then dCDN MUST only use DNS redirection and MUST include RTs to one or more surrogates in any successful RI response. CDNs MUST include the dns-only property set to True on any cascaded RI requests. Defaults to False.

An RI request for DNS-based redirection MUST include a dns dictionary. This dns dictionary MUST contain the following keys: resolver-ip, qtype, qclass, qname and the value of each MUST be the value of the appropriate part of the User Agent's DNS query/request. For internationalized domain names containing non-ASCII characters,

the value of the qname field MUST be the ASCII-compatible encoded (ACE) representation (A-label) of the domain name [RFC5890].

An example RI request (uCDN->dCDN) for DNS based redirection:

```
POST /dcdn/ri HTTP/1.1
Host: rrl.dcdn.example.net
Content-Type: application/cdni; ptype=redirection-request
Accept: application/cdni; ptype=redirection-response
```

```
{
  "dns" : {
    "resolver-ip" : "192.0.2.1",
    "c-subnet" : "198.51.100.0/24",
    "qtype" : "A",
    "qclass" : "IN",
    "qname" : "www.example.com"
  },
  "cdn-path": ["AS64496:0"],
  "max-hops": 3
}
```

4.4.2. DNS Redirection responses

For a successful DNS based redirection, the dCDN needs to return one of the following to the uCDN in the RI response:

- o The IP address(es) of (or the CNAME of) RTs that are dCDN surrogates (if the dCDN is performing DNS based redirection directly to a surrogate); or
- o The IP address(es) of (or the CNAME of) RTs that are Request Routers (if the dCDN will perform request redirection itself). A dCDN MUST NOT return a RT which is a Request Router if the dns-only key is set to True in the RI request.

The information above is encoded as a set of key:value pairs within the dns dictionary as follows:

Key	Value	Mandatory	Description
rcode	Integer	Yes	DNS response code (see [RFC6895]).
name	String	Yes	The fully qualified domain name the response relates to.
a	List of String	No	Set of IPv4 Addresses of RT(s).
aaaa	List of String	No	Set of IPv6 Addresses of RT(s).
cname	List of String	No	Set of fully qualified domain names of RT(s).
ttl	Integer	No	TTL in seconds of DNS response. Default is 0.

A successful RI response for DNS-based redirection MUST include a dns dictionary and MAY include an error dictionary (see Section 4.7). An unsuccessful RI response for DNS-based redirection MUST include an error dictionary. If a dns dictionary is included in the RI response, it MUST include the rcode and name keys and it MUST include at least one of the following keys: a, aaaa, cname. The dns dictionary MAY include both 'a' and 'aaaa' keys. If the dns dictionary contains a cname key it MUST NOT contain either an a or aaaa key. For internationalized domain names containing non-ASCII characters, the value of the cname field MUST be the ASCII-compatible encoded (ACE) representation (A-label) of the domain name.

An example of a successful RI response (dCDN->uCDN) for DNS based redirection with both a and aaaa keys is listed below :

```
HTTP/1.1 200 OK
Date: Mon, 06 Aug 2012 18:41:38 GMT
Content-Type: application/cdni; ptype=redirection-response

{
  "dns" : {
    "rcode" : 0,
    "name" : "www.example.com",
    "a" : ["203.0.113.200", "203.0.113.201", "203.0.113.202"],
    "aaaa" : ["2001:DB8::C8", "2001:DB8::C9"],
    "ttl" : 60
  }
}
```

A further example of a successful RI response (dCDN->uCDN) for DNS based redirection is listed below, in this case with a cname key containing the FQDN of the RT.

HTTP/1.1 200 OK

Date: Mon, 06 Aug 2012 18:41:38 GMT

Content-Type: application/cdni; ptype=redirection-response

```
{
  "dns" : {
    "rcode" : 0,
    "name" : "www.example.com",
    "cname" : ["rrl.dcdn.example"],
    "ttl" : 20
  }
}
```

4.5. HTTP Redirection

The following sections provide detailed descriptions of the information that should be passed in RI requests and responses for HTTP redirection.

The dictionary keys used in HTTP Redirection requests and responses use the following conventions for their prefixes:

- o c- is prefixed to keys for information related to the Client (User Agent).
- o cs- is prefixed to keys for information passed by the Client (User Agent) to the Server (uCDN).
- o sc- is prefixed to keys for information to be passed by the Server (uCDN) to the Client (User Agent).

4.5.1. HTTP Redirection requests

For HTTP-based redirection the uCDN needs to pass the following information to the dCDN in the RI request:

- o The IP address of the User Agent.
- o The URI requested by the User Agent.
- o The HTTP method requested by the User Agent
- o The HTTP version number requested by the User Agent.

The uCDN may also decide to pass the presence and value of particular HTTP headers included in the User Agent request to the dCDN.

The information above is encoded as a set of key:value pairs within the http dictionary as follows:

Key	Value	Mandatory	Description
c-ip	String	Yes	The IP address of the UA.
cs-uri	String	Yes	The Effective Request URI [RFC7230] requested by the UA.
cs-method	String	Yes	The method part of the request-line as defined in Section 3.1.1 of [RFC7230].
cs-version	String	Yes	The HTTP-version part of the request-line as defined in Section 3.1.1 of [RFC7230].
cs-(<headername>)	String	No	The field-value of the HTTP header field named <HeaderName> as a string, for example, cs-(cookie) would contain the value of the HTTP Cookie header from the UA request.

An RI request for HTTP-based redirection MUST include an http dictionary. This http dictionary MUST contain the following keys: c-ip, cs-method, cs-version and cs-uri and the value of each MUST be the value of the appropriate part of the User Agent's HTTP request.

The http dictionary of an RI request MUST contain a maximum of one cs-(<headername>) key for each unique header field-name (HTTP header field). <headername> MUST be identical to the equivalent HTTP header field-name encoded in all lowercase.

In the case where the User Agent request includes multiple HTTP header fields with the same field-name, it is RECOMMENDED that the uCDN combines these different HTTP headers into a single value according to Section 3.2.2 of [RFC7230]. However, because of the plurality of already defined HTTP header fields, and inconsistency of some of these header fields concerning the combination mechanism

defined in RFC 7230, the uCDN MAY have to deviate from using the combination mechanism where appropriate. For example, it might only send the contents of the first occurrence of the HTTP Headers instead.

An example RI request (uCDN->dCDN) for HTTP based redirection:

```
POST /dcdn/rrri HTTP/1.1
Host: rr1.dcdn.example.net
Content-Type: application/cdni; ptype=redirection-request
Accept: application/cdni; ptype=redirection-response
```

```
{
  "http": {
    "c-ip": "198.51.100.1",
    "cs-uri": "http://www.example.com",
    "cs-version": "HTTP/1.1",
    "cs-method": "GET"
  },
  "cdn-path": ["AS64496:0"],
  "max-hops": 3
}
```

4.5.2. HTTP Redirection responses

For a successful HTTP based redirection, the dCDN needs to return one of the following to the uCDN in the RI response:

- o A URI pointing to an RT that is the selected dCDN surrogate(s) (if the dCDN is performing HTTP based redirection directly to a surrogate); or
- o A URI pointing to an RT that is a Request Router (if the dCDN will perform request redirection itself).

The information above is encoded as a set of key:value pairs within the http dictionary as follows:

Key	Value	Mandatory	Description
sc-status	Integer	Yes	The status-code part of the status-line as defined in Section 3.1.2 of [RFC7230] to return to the UA (usually set to 302).
sc-version	String	Yes	The HTTP-version part of the status-line as defined in Section 3.1.2 of [RFC7230] to return to the UA.
sc-reason	String	Yes	The reason-phrase part of the status-line as defined in Section 3.1.2 of [RFC7230] to return to the UA.
cs-uri	String	Yes	The URI requested by the UA/client.
sc-(location)	String	Yes	The contents of the Location header to return to the UA (i.e., a URI pointing to the RT(s)).
sc-(<headername>)	String	No	The field-value of the HTTP header field named <HeaderName> to return to the UA. For example, sc-(expires) would contain the value of the HTTP Expires header.

Note: The sc-(location) key in the table above is an example of sc-(<headername>) that has been called out separately as its presence is mandatory in RI responses.

A successful RI response for HTTP-based redirection MUST include an http dictionary and MAY include an error dictionary (see Section 4.7). An unsuccessful RI response for HTTP-based redirection MUST include an error dictionary. If an http dictionary is included in the RI response, it MUST include at least the following keys: sc-status, sc-version, sc-reason, cs-uri and sc-(location).

The http dictionary of an RI response MUST contain a maximum of one sc-(<headername>) key for each unique header field-name (HTTP header field). <headername> MUST be identical to the equivalent HTTP header field-name encoded in all lowercase.

The uCDN MAY decide to not return, override or alter any or all of the HTTP headers defined by sc-(<headername>) keys before sending the HTTP response to the UA. It should be noted that in some cases, sending the HTTP Headers indicated by the dCDN transparently on to the UA might result in, for the uCDN, undesired behaviour. As an example, the dCDN might include sc-(cache-control), sc-(last-modified) and sc-(expires) keys in the http dictionary, through which the dCDN may try to influence the cacheability of the response by the UA. If the uCDN would pass these HTTP headers on to the UA, this could mean that further requests from the uCDN would go directly to the dCDN, bypassing the uCDN and any logging it may perform on incoming requests. The uCDN is therefore recommended to carefully consider which HTTP headers to pass on, and which to either override or not pass on at all.

An example of a successful RI response (dCDN->uCDN) for HTTP based redirection:

```
HTTP/1.1 200 OK
Date: Mon, 06 Aug 2012 18:41:38 GMT
Content-Type: application/cdni; ptype=redirection-response

{
  "http": {
    "sc-status": 302,
    "sc-version": "HTTP/1.1",
    "sc-reason": "Found",
    "cs-uri": "http://www.example.com"
    "sc-(location)":
      "http://surl.dcdn.example/ucdn/example.com",
  }
}
```

4.6. Cacheability and scope of responses

RI responses may be cacheable. As long as a cached RI response is not stale according to standard HTTP Cache-Control or other applicable mechanisms, it may be reused by the uCDN in response to User Agent requests without sending another RI request to the dCDN.

An RI response MUST NOT be reused unless the request from the User Agent would generate an identical RI request to the dCDN as the one that resulted in the cached RI response (except for the c-ip field

provided that the User Agent's c-ip is covered by the scope in the original RI response, as elaborated upon below).

Additionally, although RI requests only encode a single User Agent request to be redirected there may be cases where a dCDN wishes to indicate to the uCDN that the RI response can be reused for other User Agent requests without the uCDN having to make another request via the RI. For example, a dCDN may know that it will always select the same Surrogates for a given set of User Agent IP addresses and in order to reduce request volume across the RI or to remove the additional latency associated with an RI request, the dCDN may wish to indicate that set of User Agent IP addresses to the uCDN in the initial RI response. This is achieved by including an optional scope dictionary in the RI response.

Scope is encoded as a set of key:value pairs within the scope dictionary as follows:

Key	Value	Mandatory	Description
iprange	List of String	No	A List of IP subnets in CIDR notation that this RI response can be reused for, provided the RI response is still considered fresh.

If a uCDN has multiple cached responses with overlapping scopes and a UA request comes in for which the User Agent's IP matches with the IP subnets in multiple of these cached responses, the uCDN SHOULD use the most recent cached response when determining the appropriate RI response to use.

The following is an example of a DNS redirection response from Section 4.4.2 that is cacheable by the uCDN for 30 seconds and can be returned to any User Agent with an IPv4 address in 198.51.100.0/24.

```
HTTP/1.1 200 OK
Date: Mon, 06 Aug 2012 18:41:38 GMT
Content-Type: application/cdni; ptype=redirection-response
Cache-Control: public, max-age=30
```

```
{
  "dns" : {
    "rcode" : 0,
    "name" : "www.example.com",
    "a" : ["203.0.113.200", "203.0.113.201"],
    "aaaa" : ["2001:DB8::C8", "2001:DB8::C9"],
    "ttl" : 60
  }
  "scope" : {
    "iprange" : ["198.51.100.0/24"]
  }
}
```

Example of HTTP redirection response from Section 4.5.2 that is cacheable by the uCDN for 60 seconds and can be returned to any User Agent with an IPv4 address in 198.51.100.0/24.

Note: The response to the UA is only valid for 30 seconds, whereas the uCDN can cache the RI response for 60 seconds.

```
HTTP/1.1 200 OK
Date: Mon, 06 Aug 2012 18:41:38 GMT
Content-Type: application/cdni; ptype=redirection-response
Cache-Control: public, max-age=60
```

```
{
  "http": {
    "sc-status": 302,
    "cs-uri": "http://www.example.com"
    "sc-(location)":
      "http://surl.dcdn.example/ucdn/example.com",
    "sc-(cache-control)" : "public, max-age=30"
  }
  "scope" : {
    "iprange" : ["198.51.100.0/24"]
  }
}
```

4.7. Error responses

From a uCDN perspective, there are two types of errors that can be the result of the transmission of an RI request to a dCDN:

1. An HTTP protocol error signaled via an HTTP status code, indicating a problem with the reception or parsing of the RI request or the generation of the RI response by the dCDN, and
2. An RI-level error specified in an RI response message

This section deals with the latter type. The former type is outside the scope of this document.

There are numerous reasons for a dCDN to be unable to return an affirmative RI response to a uCDN. Reasons may include both dCDN internal issues such as capacity problems, as well as reasons outside the influence of the dCDN, such as a malformed RI request. To aid with diagnosing the cause of errors, RI responses SHOULD include an error dictionary to provide additional information to the uCDN as to the reason/cause of the error. The intention behind the error dictionary is to aid with either manual or automatic diagnosis of issues. The resolution of such issues is outside the scope of this document; this document does not specify any consequent actions a uCDN should take upon receiving a particular error code.

Error information (if present) is encoded as a set of key:value pairs within a JSON-encoded error dictionary as follows:

Key	Value	Mandatory	Description
error-code	Integer	Yes	A three-digit numeric code defined by the server to indicate the error(s) that occurred.
reason	String	No	A string providing further information related to the error.

The first digit of the error-code defines the class of error. There are 5 classes of error distinguished by the first digit of the error-code:

1xx: Informational (no error): The response should not be considered an error by the uCDN, which may proceed by redirecting the UA according to the values in the RI response. The error code and accompanying description may be used for informational purposes, e.g., for logging.

2xx: Reserved.

3xx: Reserved.

4xx: uCDN error: The dCDN can not or will not process the request due to something that is perceived to be a uCDN error, for example, the RI request could not be parsed successfully by the dCDN. The last two-digits may be used to more specifically indicate the source of the problem.

5xx: dCDN error: Indicates that the dCDN is aware that it has erred or is incapable of satisfying the RI request for some reason, for example, the dCDN was able to parse the RI request but encountered an error for some reason. Examples include the dCDN not being able to retrieve the associated metadata or the dCDN being out of capacity.

The following error codes are defined and maintained by IANA (see Section 6):

Error codes with a "Reason" of "<reason>" do not have a defined value for their 'reason'-key. Depending on the error-code semantics, the value of this field may be determined dynamically.

Code	Reason	Description
100	<reason> (see Description)	Generic informational error-code meant for carrying a human-readable string
400	<reason> (see Description)	Generic error-code for uCDN errors where the dCDN can not or will not process the request due to something that is perceived to be a uCDN error. The reason field may be used to provide more details about the source of the error.
500	<reason> (see Description)	Generic error-code for dCDN errors where the dCDN is aware that it has erred or is incapable of satisfying the RI request for some reason. The reason field may be used to provide more details about the source of the error.
501	Unable to retrieve metadata	The dCDN is unable to retrieve the metadata associated with the content requested by the UA. This may indicate a configuration error or the content requested by the UA not existing.
502	Loop detected	The dCDN detected a redirection loop (see Section 4.8).
503	Maximum hops exceeded	The dCDN detected the maximum number of redirection hops exceeding max-hops (see Section 4.8).
504	Out of capacity	The dCDN does not currently have sufficient capacity to handle the UA request.
505	Delivery protocol not supported	The dCDN does not support the (set of) delivery protocols indicated in the CDNI Metadata of the content requested content by the UA.
506	Redirection protocol not supported	The dCDN does not support the requested redirection protocol. This error-code is also used when the RI request has the dns-only flag set to True and the dCDN is not support or is not prepared to return a RT of a surrogate directly.

Table 1

The following is an example of an unsuccessful RI response (dCDN->uCDN) for a DNS based User Agent request:

```
HTTP/1.1 500 Internal Server Error
Date: Mon, 06 Aug 2012 18:41:38 GMT
Content-Type: application/cdni; ptype=redirection-response
Cache-Control: private, no-cache
```

```
{
  "error" : {
    "error-code" : 504,
    "description" : "Out of capacity"
  }
}
```

The following is an example of a successful RI response (dCDN->uCDN) for a HTTP based User Agent request containing an error dictionary for informational purposes:

```
HTTP/1.1 200 OK
Date: Mon, 06 Aug 2012 18:41:38 GMT
Content-Type: application/cdni; ptype=redirection-response
Cache-Control: private, no-cache
```

```
{
  "http": {
    "sc-status": 302,
    "sc-version": "HTTP/1.1",
    "sc-reason": "Found",
    "cs-uri": "http://www.example.com"
    "sc-(location)":
      "http://surl.dcdn.example/ucdn/example.com",
  },
  "error" : {
    "error-code" : 100,
    "description" :
      "This is a human-readable message meant for debugging purposes"
  }
}
```

4.8. Loop detection & prevention

In order to prevent and detect RI request loops, each CDN MUST insert its CDN Provider ID into the cdn-path key of every RI request it originates or cascades. When receiving RI requests a dCDN MUST check the cdn-path and reject any RI requests which already contain the dCDN's Provider ID in the cdn-path. Transit CDNs MUST NOT propagate to any downstream CDNs if the number of CDN Provider IDs in cdn-path (before adding its own Provider ID) is equal to or greater than max-hops.

The CDN Provider ID uniquely identifies each CDN provider during the course of request routing redirection. It consists of the characters AS followed by the CDN Provider's AS number, then a colon (':') and an additional qualifier that is used to guarantee uniqueness in case a particular AS has multiple independent CDNs deployed. For example, "AS64496:0".

If a dCDN receives an RI request whose cdn-path already contains that dCDN's Provider ID the dCDN MUST send an RI error response which SHOULD include an error code of 502.

If a dCDN receives an RI request where the number of CDN Provider IDs in cdn-path is greater than max-hops, the dCDN MUST send an RI error response which SHOULD include an error code of 503.

It should be noted that the loop detection & prevention mechanisms described above only cover preventing and detecting loops within the RI itself. Besides loops within the RI itself, there is also the possibility of loops in the data plane, for example, if the IP address(es) or URI(s) returned in RI responses do not resolve directly to a surrogate in the final dCDN there is the possibility that a User Agent may be continuously redirected through a loop of CDNs. The specification of solutions to address data plane request redirection loops between CDNs is outside of the scope of this document.

5. Security Considerations

Information passed over the RI could be considered personal or sensitive, for example, RI requests contain parts of a User Agent's original request and RI responses reveal information about the dCDN's policy for which surrogates should serve which content/user locations.

The RI interface also provides a mechanism whereby a uCDN could probe a dCDN and infer the dCDN's edge topology by making repeated RI requests for different content and/or UA IP addresses and correlating the responses from the dCDN. Additionally the ability for a dCDN to indicate that an RI response applies more widely than the original request (via the scope dictionary) may significantly reduce the number of RI requests required to probe and infer the dCDN's edge topology.

The same information could be obtained in the absence of the RI interface, but it could be more difficult to gather as it would require a distributed set of machines with a range of different IP addresses each making requests directly to the dCDN. However, the RI facilitates easier collection of such information as it enables a

single client to query the dCDN for a redirection/surrogate selection on behalf of any UA IP address.

5.1. Authentication, Authorization, Confidentiality, Integrity Protection

An implementation of the CDNI Redirection interface MUST support TLS transport as per [RFC2818] and [RFC7230]. The use of TLS for transport of the CDNI Redirection interface messages allows:

- o The dCDN and uCDN to authenticate each other

and, once they have mutually authenticated each other, it allows:

- o The dCDN and uCDN to authorize each other (to ensure they are transmitting/receiving CDNI Redirection messages to/from an authorized CDN);
- o CDNI Redirection interface messages to be transmitted with confidentiality; and
- o The integrity of the CDNI Redirection interface messages to be protected during the exchange.

In an environment where any such protection is required, mutually authenticated encrypted transport MUST be used to ensure confidentiality of the redirection information, and to do so, TLS MUST be used (including authentication of the remote end) by the server-side (dCDN) and the client-side (uCDN) of the CDNI Redirection interface.

When TLS is used, the general TLS usage guidance in [RFC7525] MUST be followed.

5.2. Privacy

Information passed over the RI ought to be considered personal and sensitive. In particular, parts of a User Agent's original request, most notably the UA's IP address and requested URI, are transmitted over the RI to the dCDN. The use of mutually authenticated TLS, as described in the previous section, prevents any other party than the authorized dCDN from gaining access to this information.

Regardless of whether the uCDN and dCDN use the RI, a successful redirect from a uCDN to a dCDN will make that dCDN aware of the UA's IP address. As such, the fact that this information is transmitted across the RI does not allow the dCDN to learn new information. On the other hand, if a uCDN uses the RI to check with multiple

candidate dCDNs, those candidates that do not end up getting redirected to, do obtain information regarding End User IP addresses and requested URIs that they would not have, had the RI not been used.

While it is technically possible to mask some information in the RI Request, such as the last bits of the UA IP address, it is important to note that this will reduce the effectiveness of the RI in certain cases. CDN deployments need to strike a balance between end-user privacy and the features impacted by such masking. This balance is likely to vary from one deployment to another. As an example, when the UA and its DNS resolver is behind a Carrier-grade NAT, and the RI is used to find an appropriate delivery node behind the same NAT, the full IP address might be necessary. Another potential issue when using IP anonymization is that it is no longer possible to correlate an RI Request with a subsequent UA request.

6. IANA Considerations

6.1. CDNI Payload Type Parameter registrations

The IANA is requested to register the following two new Payload Types in the CDNI Payload Type Parameter registry for use with the application/cdni MIME media type.

[RFC Editor Note: Please replace the references to [RFCthis] below with this document's RFC number before publication.]

Payload Type	Specification
redirection-request	[RFCthis]
redirection-response	[RFCthis]

6.1.1. CDNI RI Redirection Request Payload Type

Purpose: The purpose of this payload type is to distinguish RI request messages.

Interface: RI

Encoding: see Section 4.4.1 and Section 4.5.1

6.1.2. CDNI RI Redirection Response Payload Type

Purpose: The purpose of this payload type is to distinguish RI response messages.

Interface: RI

Encoding: see Section 4.4.2 and Section 4.5.2

6.2. RI Error response registry

IANA is requested to create a new "CDNI RI Error response code" subregistry within the "Content Delivery Network Interconnection (CDNI) Parameters" registry. The "CDNI RI Error response code" namespace defines the valid values for the error-code key in RI error responses. The CDNI RI Error response code MUST be a three digit integer.

Additions to the "RI Error response registry" will be made via "Specification Required" as defined in [RFC5226].

The Designated Expert will verify that new error code registrations do not duplicate existing error code definitions (in name or functionality), ensure that the new error code is in accordance with the error classes defined in section Section 4.7 of this document, prevent gratuitous additions to the namespace, and prevent any additions to the namespace that would impair the interoperability of CDNI implementations.

New registrations are required to provide the following information:

Code: A three-digit numeric error-code, in accordance with the error classes defined in section Section 4.7 of this document.

Reason: A string that provides further information related to the error that will be included in the JSON error dictionary with the 'reason'-key. Depending on the error-code semantics, the value of this field may be determined dynamically. In that case, the registration should set this value to '<reason>' and define its semantics in the description field.

Description: A brief description of the error code semantics.

Specification: Reference to the specification that defines the error code in more detail.

The entries in Table 1 are registered by this document, with the value of the 'Specification' field set to [RFCThis].

7. Contributors

[RFC Editor Note: Please move the contents of this section to the Authors' Addresses section prior to publication as an RFC.]

The following persons have participated as co-authors to this document:

Wang Danhua, Huawei, Email: wangdanhua@huawei.com

He Xiaoyan, Huawei, Email: hexiaoyan@huawei.com

Ge Chen, China Telecom, Email: cheng@gsta.com

Ni Wei, China Mobile, Email: niwei@chinamobile.com

Yunfei Zhang, Email: hishigh@gmail.com

Spencer Dawkins, Huawei, Email: spencer@wonderhamster.org

8. Acknowledgements

The authors would like to thank Taesang Choi, Francois le Faucheur, Matt Miller, Scott Wainner and Kevin J Ma for their valuable comments and input to this document.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, DOI 10.17487/RFC3986, January 2005, <<http://www.rfc-editor.org/info/rfc3986>>.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<http://www.rfc-editor.org/info/rfc4291>>.
- [RFC5952] Kawamura, S. and M. Kawashima, "A Recommendation for IPv6 Address Text Representation", RFC 5952, DOI 10.17487/RFC5952, August 2010, <<http://www.rfc-editor.org/info/rfc5952>>.

- [RFC7230] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", RFC 7230, DOI 10.17487/RFC7230, June 2014, <<http://www.rfc-editor.org/info/rfc7230>>.
- [RFC7159] Bray, T., Ed., "The JavaScript Object Notation (JSON) Data Interchange Format", RFC 7159, DOI 10.17487/RFC7159, March 2014, <<http://www.rfc-editor.org/info/rfc7159>>.
- [RFC6895] Eastlake 3rd, D., "Domain Name System (DNS) IANA Considerations", BCP 42, RFC 6895, DOI 10.17487/RFC6895, April 2013, <<http://www.rfc-editor.org/info/rfc6895>>.
- [RFC7493] Bray, T., Ed., "The I-JSON Message Format", RFC 7493, DOI 10.17487/RFC7493, March 2015, <<http://www.rfc-editor.org/info/rfc7493>>.
- [RFC7525] Sheffer, Y., Holz, R., and P. Saint-Andre, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", BCP 195, RFC 7525, DOI 10.17487/RFC7525, May 2015, <<http://www.rfc-editor.org/info/rfc7525>>.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, DOI 10.17487/RFC5226, May 2008, <<http://www.rfc-editor.org/info/rfc5226>>.
- [RFC6707] Niven-Jenkins, B., Le Faucheur, F., and N. Bitar, "Content Distribution Network Interconnection (CDNI) Problem Statement", RFC 6707, DOI 10.17487/RFC6707, September 2012, <<http://www.rfc-editor.org/info/rfc6707>>.
- [RFC2818] Rescorla, E., "HTTP Over TLS", RFC 2818, DOI 10.17487/RFC2818, May 2000, <<http://www.rfc-editor.org/info/rfc2818>>.
- [RFC5890] Klensin, J., "Internationalized Domain Names for Applications (IDNA): Definitions and Document Framework", RFC 5890, DOI 10.17487/RFC5890, August 2010, <<http://www.rfc-editor.org/info/rfc5890>>.
- [RTMP] Adobe Systems Incorporated, "Real-Time Messaging Protocol (RTMP) specification", December 2012, <http://www.adobe.com/go/spec_rtmp>.

9.2. Informative References

- [RFC7337] Leung, K., Ed. and Y. Lee, Ed., "Content Distribution Network Interconnection (CDNI) Requirements", RFC 7337, DOI 10.17487/RFC7337, August 2014, <<http://www.rfc-editor.org/info/rfc7337>>.
- [RFC7336] Peterson, L., Davie, B., and R. van Brandenburg, Ed., "Framework for Content Distribution Network Interconnection (CDNI)", RFC 7336, DOI 10.17487/RFC7336, August 2014, <<http://www.rfc-editor.org/info/rfc7336>>.
- [RFC7736] Ma, K., "Content Delivery Network Interconnection (CDNI) Media Type Registration", RFC 7736, DOI 10.17487/RFC7736, December 2015, <<http://www.rfc-editor.org/info/rfc7736>>.

Authors' Addresses

Ben Niven-Jenkins (editor)
Nokia
3 Ely Road
Milton, Cambridge CB24 6DD
UK

Email: ben.niven-jenkins@nokia.com

Ray van Brandenburg (editor)
TNO
Anna van Buerenplein 1
The Hague 2595DA
the Netherlands

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

CDNI
Internet-Draft
Intended status: Standards Track
Expires: March 3, 2014

K. Leung
F. Le Faucheur
Cisco Systems
B. Downey
Verizon Labs
R. van Brandenburg
TNO
S. Leibrand
Limelight Networks
August 30, 2013

URI Signing for CDN Interconnection (CDNI)
draft-leung-cdni-uri-signing-03

Abstract

This document describes how the concept of URI signing supports the content access control requirements of CDNI and proposes a candidate URI signing scheme.

The proposed URI signing method specifies the information needed to be included in the URI and the algorithm used to authorize and to validate access request for the content referenced by the URI. Some of the information may be accessed by the CDN via configuration or CDNI metadata.

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 March 3, 2014.

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	3
1.1. Terminology	3
1.2. Background on URI Signing	4
1.3. CDNI URI Signing Overview	6
2. Signed URI Format	8
2.1. Enforcement Attributes	9
2.2. Signature Computation Attributes	10
2.3. URI Signature Attributes	11
2.4. URI Signing Package Attribute	11
3. Signing a URI	12
4. Validating a URI Signature	17
5. Considerations for CDNI Interfaces	19
5.1. CDNI Control Interface	20
5.2. CDNI Footprint & Capabilities Advertisement Interface	20
5.3. CDNI Request Routing Redirection Interface	20
5.4. CDNI Metadata Interface	21
5.5. CDNI Logging Interface	22
6. Detailed URI Signing Operation	22
6.1. HTTP Redirection	23
6.2. DNS Redirection	25
7. HTTP Adaptive Streaming	28
8. IANA Considerations	28
9. Security Considerations	29
10. Privacy	31
11. Acknowledgements	31
12. References	31
12.1. Normative References	31
12.2. Informative References	31
Authors' Addresses	32

1. Introduction

This document describes the concept of URI Signing and how it can be used to provide access authorization in the case of interconnected CDNs (CDNI). The primary goal of URI Signing is to make sure that only authorized User Agents (UAs) are able to access the content, with a Content Service Provider (CSP) being able to authorize every individual request. It should be noted that URI Signing is not a content protection scheme; if a CSP wants to protect the content itself, other mechanisms, such as DRM, are more appropriate.

The overall problem space for CDN Interconnection (CDNI) is described in CDNI Problem Statement [RFC6707]. In this document, along with the CDNI Requirements [I-D.ietf-cdni-requirements] document and the CDNI Framework [I-D.ietf-cdni-framework] the need for interconnected CDNs to be able to implement an access control mechanism that enforces the CSP's distribution policy is described.

Specifically, CDNI Framework [I-D.ietf-cdni-framework] states:

"The CSP may also trust the CDN operator to perform actions such as ..., and to enforce per-request authorization performed by the CSP using techniques such as URI signing."

In particular, the following requirement is listed in CDNI Requirements [I-D.ietf-cdni-requirements]:

"MI-16 [HIGH] The CDNI Metadata Distribution interface shall allow signaling of authorization checks and validation that are to be performed by the surrogate before delivery. For example, this could potentially include:

- * need to validate URI signed information (e.g. Expiry time, Client IP address)."

This document proposes a URI Signing scheme that allows Surrogates in interconnected CDNs to enforce a per-request authorization performed by the CSP. Splitting the role of performing per-request authorization by CSP and the role of validation of this authorization by the CDN allows any arbitrary distribution policy to be enforced across CDNs without the need of CDNs to have any awareness of the actual CSP distribution policy.

1.1. Terminology

This document uses the terminology defined in CDNI Problem Statement [RFC6707].

This document also uses the terminology of Keyed-Hashing for Message Authentication (HMAC) [RFC2104] including the following terms (reproduced here for convenience):

- o MAC: message authentication code.
- o HMAC: Hash-based message authentication code (HMAC) is a specific construction for calculating a MAC involving a cryptographic hash function in combination with a secret key.
- o HMAC-SHA1: HMAC instantiation using SHA1 as the cryptographic hash function.
- o HMAC-MD5: HMAC instantiation using MD5 as the cryptographic hash function.

In addition, the following terms are used throughout this document:

- o URI Signature: Message digest or digital signature that is computed with an algorithm for protecting the URI.
- o Original URI: The URI before URI Signing is applied.
- o Signed URI: Any URI that contains a URI signature.
- o Target CDN URI: Embedded URI created by the CSP to direct UA towards the Upstream CDN. The Target CDN URI can be signed by the CSP and verified by the Upstream CDN.
- o Redirection URI: URI created by the Upstream CDN to redirect UA towards the Downstream CDN. The Redirection URI can be signed by the Upstream CDN and verified by the Downstream CDN. In a cascaded CDNI scenario, there can be more than one Redirection URI.

1.2. Background on URI Signing

The next section provides an overview of how URI Signing works in a CDNI environment. As background information, URI Signing is first explained in terms of a single CDN delivering content on behalf of a CSP.

A CSP and CDN are assumed to have a trust relationship that enables the CSP to authorize access to a content item by including a set of attributes in the URI before redirecting a UA to the CDN. Using these attributes, it is possible for a CDN to check an incoming content request to see whether it was authorized by the CSP (e.g. based on the UA's IP address or a time window). Of course, the

attributes need to be added to the URI in a way that prevents a UA from changing the attributes, thereby leaving the CDN to think that the request was authorized by the CSP when in fact it wasn't. For this reason, a URI Signing mechanism includes in the URI a message digest or digital signature that allows a CDN to check the authenticity of the URI. The message digest or digital signature can be calculated based on a shared secret between the CSP and CDN or using CSP's asymmetric public/private key pair, respectively.

Figure 1, shown below, presents an overview of the URI Signing mechanism in the case of a CSP with a single CDN. When the UA browses for content on CSP's website (#1), it receives HTML web pages with embedded content URIs. Upon requesting these URIs, the CSP redirects to a CDN, creating a Target CDN URI (#2) (alternatively, the target CDN URI itself is embedded in the HTML). The Target CDN URI is the signed URI which may include the IP address of the UA and/or a time window and always contains the URI signature which is generated by the CSP using the shared secret or a private key. Once the UA receives the response with the embedded URI, it sends a new HTTP request using the embedded URI to the CDN (#3). Upon receiving the request, the CDN checks to see if the Signed URI is authentic by verifying the URI signature. In addition, it checks whether the IP address of the HTTP request matches that in the Signed URI and if the time window is still valid. After these values are confirmed to be valid, the CDN delivers the content (#4).

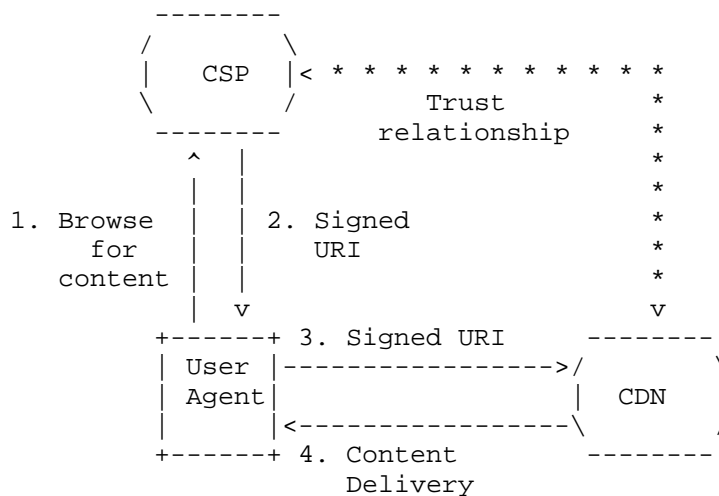


Figure 1: URI Signing in a CDN Environment

1.3. CDNI URI Signing Overview

In a CDNI environment, URI Signing operates the same way in the initial steps #1 and #2 but the later steps involve multiple CDNs in the process of delivering the content. The main difference from the single CDN case is a redirection step between the Upstream CDN and the Downstream CDN. In step #3, UA may send HTTP request or DNS request. Depending on whether HTTP-based or DNS-based request routing is used, the Upstream CDN responds by directing the UA towards the Downstream CDN using either a Redirection URI (which is a Signed URI generated by the Upstream CDN) or a DNS reply, respectively (#4). Once the UA receives the response, it sends the Redirection URI/Target CDN URI to the Downstream CDN (#5). The received URI is validated by the Downstream CDN before delivering the content (#6). This is depicted in the figure below. Note: The CDNI call flows are covered in Detailed URI Signing Operation (Section 6).

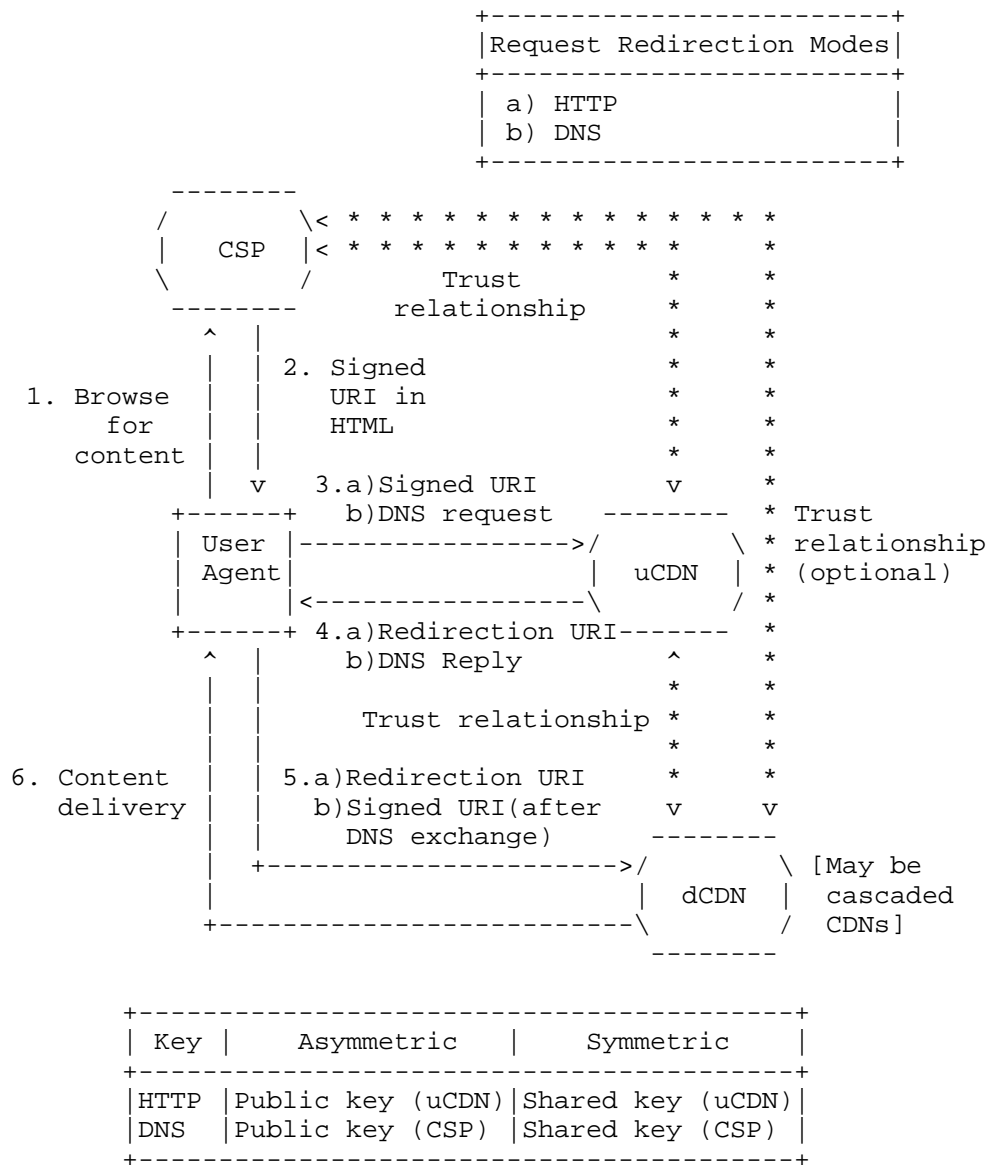


Figure 2: URI Signing in a CDNI Environment

The trust relationships between CSP, Upstream CDN, and Downstream CDN have direct implications for URI Signing. In the case shown in Figure 2, the CDN that the CSP has a trust relationship with is the Upstream CDN. The delivery of the content may be delegated to the Downstream CDN, which has a relationship with the Upstream CDN but

may have no relationship with the CSP.

In CDNI, there are two methods for request routing: DNS-based and HTTP-based. For DNS-based request routing, the Signed URI (i.e. Target CDN URI) provided by the CSP reaches the Downstream CDN directly. In the case where the Downstream CDN does not have a trust relationship with the CSP, this means that only an asymmetric public/private key method can be used for computing the URI signature because the CSP and Downstream CDN are not able to exchange symmetric shared secret keys. Since the CSP is unlikely to have relationships with all the Downstream CDNs that are delegated to by the Upstream CDN, CSP may choose to allow the Authoritative CDN to redistribute the shared key to a subset of their Downstream CDNs .

For HTTP-based request routing, the Signed URI (i.e. Target CDN URI) provided by the CSP reaches the Upstream CDN. After this URI has been verified to be correct by the Upstream CDN, the Upstream CDN creates and signs a new Redirection URI to redirect the UA to the Downstream CDN. Since this new URI also has a new URI signature, this new signature can be based around the trust relationship between the Upstream CDN and Downstream CDN, and the relationship between the Downstream CDN and CSP is not relevant. Given the fact that such a relationship between Upstream CDN and Downstream CDN always exists, both asymmetric public/private keys and symmetric shared secret keys can be used for URI Signing. Note that the signed Redirection URI needs to maintain the same level of security as the original Signed URI.

2. Signed URI Format

The concept behind URI Signing is based on embedding in the Target CDN URI/Redirection URI the CDNI attributes that can be validated to ensure the UA has legitimate access to the content. The CDNI attribute(s) are appended to the original URI.

NOTE: If the UA or another entity needs to add one or more attributes to the Signed URI for purposes other than URI Signing, those attribute MUST be appended after the CDNI URI Signing attributes discussed in this document. Any attributes appended in such way after the URI signature has been calculated are not validated for the purpose of content access authorization. Note that adding any such attributes to the Signed URI before the CDNI URI Signing attributes will cause the URI Signing validation to fail.

For the purposes of the URI signing mechanism described in this document, four types of attributes may be embedded in the URI:

- o Enforcement Attributes: Attributes that are used to enforce a distribution policy defined by the CSP. Examples of enforcement attributes are IP address of the UA and time window.
- o Signature Computation Attributes: Attributes that are used by the CDN to verify the URI signature embedded in the received URI. In order to verify a URI signature, the CDN requires some attributes that describe how the URI signature was generated. Examples of Signature Computation Attributes include the used HMAC's hash function and/or the key identifier.
- o URI Signature Attributes: The attribute that carries the actual message digest or digital signature representing the URI signature used for checking the integrity and authenticity of the URI.
- o URI Signing Package Attribute: The attribute that encapsulates all the other URI Signing attributes in an encoded format. When this attribute is used, it is the only URI Signing attribute that is exposed in the Signed URI.

Two types of keys can be used for URI Signing: asymmetric keys and symmetric keys. Asymmetric keys are based on a public/private key pair mechanism and always contain a private key only known to the CDN (or CSP) signing the URI and a public key for the verification of the Signed URI. Regardless of the type of key used, the entity that validates the URI has to obtain the key. There are very different requirements for key distribution (out of scope of this document) with asymmetric keys and with symmetric keys. Key distribution for symmetric keys requires confidentiality to prevent another party from getting access to the key, since it could then generate valid Signed URIs for unauthorized requests. Key distribution for asymmetric keys does not require confidentiality since public keys can typically be distributed openly (because they cannot be used for URI signing) and private keys are kept by the URI signing function.

2.1. Enforcement Attributes

This section identifies the set of attributes that may be needed to enforce the CSP distribution policy. These attributes are protected by the URI signature. New attributes may be introduced in the future to extend the capabilities of the distribution policy.

In order to provide flexibility in distribution policies to be enforced, the exact subset of attributes used for URI signature in a given request is a deployment decision. The defined keyword for each query string attribute is specified in parenthesis below.

The following attributes are used to enforce the distribution policy:

- o Expiry Time (ET) [optional] - Time when the Signed URI expires. This is represented in seconds since midnight 1/1/1970 UTC (i.e. UNIX epoch). The request is rejected if the received time is later than this timestamp. Note: The time including time zone on the entities that generate and validate the signed URI need to be in sync (e.g. NTP is used).
- o Client IP (CIP) [optional] - IP address of the client for which this signed URI is generated. This is represented in dotted decimal format for IPv4 or canonical text representation for IPv6 address [RFC5952] . The request is rejected if sourced from a client with a different IP address.

The Expiry Time attribute ensures that the content authorization expires after a predetermined time. This limits the time window for content access and prevents replay of the request beyond the authorized time window.

The Client IP attribute is used to restrict content access to a particular User Agent, based on its IP address for whom the content access was authorized.

2.2. Signature Computation Attributes

This section identifies the set of attributes that may be needed to verify the URI (signature). New attributes may be introduced in the future if new URI signing algorithms are developed.

The defined keyword for each query string attribute is specified in parenthesis below.

The following attributes are used to verify the URI (signature).

- o Version (VER) [optional] - An integer used for identifying the version of URI signing method.
- o Key ID (KID) [optional] - A string used for obtaining the key (e.g. database lookup, URI reference) which is needed to validate the URI signature.
- o Hash Function (HF) [optional] - A string used for identifying the hash function to compute the URI signature (e.g. "MD5", "SHA1") with HMAC. If this attribute is not present in the URI, the default hash function is SHA256.

The Version attribute indicates which version of URI signing scheme is used (including which attributes and algorithms are supported). The present document specifies Version 1. If the Version attribute

is not present in the Signed URI, then the version is obtained from the CDNI metadata, else it is considered to have been set to the default value of 1. More versions may be defined in the future.

The Key ID attribute is used to retrieve the key which is needed as input to the algorithm for validating the Signed URI.

The Hash Function attribute indicates the hash function to be used for HMAC-based message digest computation.

2.3. URI Signature Attributes

The following attributes are used to convey the actual URI signature. Just to reiterate, the protected portion of the Signed URI is the entire URI, excluding the scheme name part, that is "in front" of either of the attributes below. The attribute MUST be the last attribute in the query string of the URI after all the CDNI URI Signing attributes are appended to the URI.

- o Message Digest (MD) [mandatory for symmetric key] - A string used for the message digest generated by the URI signer.
- o Digital Signature (DS) [mandatory for asymmetric keys] - A string used for the digital signature provided by the URI signer.

The Message Digest attribute contains the message digest used to validate the Signed URI when symmetric key is used. In the case of symmetric key, HMAC algorithm is used for the following reasons: 1) Ability to use hash functions (i.e. no changes needed) with well understood cryptographic properties that perform well and for which code is freely and widely available, 2) Easy to replace the embedded hash function in case faster or more secure hash functions are found or required, 3) Original performance of the hash function is maintained without incurring a significant degradation, and 4) Simple way to use and handle keys.

The Digital Signature attribute contains the digital signature used to verify the Signed URI when asymmetric keys are used. In the case of asymmetric keys, Elliptic Curve Digital Signature Algorithm (ECDSA) - a variant of DSA - is used because of the following reasons: 1) Key size is small while still offering good security, 2) Key is easy to store, and 3) Computation is faster than DSA or RSA. [Editor's note: Should there be options for other DSA?]

2.4. URI Signing Package Attribute

The URI Signing Package attribute is a container for all the other CDNI URI Signing attributes in an encoded format. The CDNI URI

Signing attributes used to compute the URI signature are encoded and stored in this attribute. It is appended to the original URI to create the Signed URI. The URI Signing Package attribute serves multiple functions. It avoids the exposure of all the other CDNI URI Signing attributes. For those attributes, the potential for attribute namespace conflict is eliminated. Overlapping attribute names cause ambiguity for the optional CDNI URI Signing attributes because it's unclear if the attributes in the query string were added for URI Signing or another purpose. One way to prevent the situation is to require the CSP/CDN to not use the same attribute names when generating an URI. But that may not be always possible. Another benefit of the attribute is that the obfuscation hides the information (e.g. client IP address) from view for the common user, whom is not aware of the encoding scheme. Obviously, it is not a security method since anyone who knows the encoding scheme is able to obtain the clear text. Note that any attributes that are appended after the URI Signing Package attribute are not validated and hence do not affect URI Signing. The attributes that are appended to the Signed URI should not have the same name as this attribute to avoid failing URI Signing validation.[Editor's note: discuss if URI Signing Package should be mandatory]

The following attribute is used to carry the encoded set of URI Signing attributes in the Signed URI.

- o URI Signing Package (URISigningToken) - The encoded token containing all the CDNI URI Signing attributes used for URI Signing.

The URI Signing Package attribute contains the URI Signing attributes in the Base 64 encoding with URL and Filename Safe Alphabet (a.k.a. "base64url") as specified in the Base-64 Data Encoding [RFC4648] document. When this attribute is used, it is the only URI Signing attribute exposed in the Signed URI. The attribute MUST be the last attribute in the query string of the URI when the Signed URI is generated. However, a client or CDN may append other attributes unrelated to URI Signing to the Signed URI as long as those attribute names do not use the same name as this attribute. The CDNI Metadata Interface may override the attribute name (i.e. "URISigningToken") and/or the encoding format used in the URI Signing Package attribute.

3. Signing a URI

The following procedure for signing a URI defines the algorithms in this version of URI Signing. Note that some steps may be skipped if the URI Signing attribute is not needed to enforce the distribution policy. A URI (as defined in URI Generic Syntax [RFC3986]) contains

the following parts: scheme name, authority, path, query, and fragment. The entire URI except the "scheme name" part is protected by the URI signature. This allows the URI signature to be validated correctly in the case when a client performs a fallback to another scheme (e.g. HTTP) for a content referenced by an URI with a specific scheme (e.g. RTSP). The benefit is that the content access is protected regardless of the type of transport used for delivery. If the CSP wants to ensure a specific protocol is used for content delivery, that information is passed by CDNI metadata.

Note: The following URI signing steps are specified to generate a Signed URI. However, it is possible to use some other algorithm and implementation as long as the same result is achieved. An example for the Original URI, "http://example.com/content.mov", is used to clarify the steps.

The URI Signing attributes are appended to the protected portion of the URI to compute the URI signature.

1. Copy the Original URI, excluding the "scheme name" part, into a buffer to hold the message for performing the operations below.
2. Check if the URI already contains a query string. If not, append a "?" character. If yes, append an "&" character.
3. If the version needs to be specified, then append the string "VER=1". This represents the version of URI Signing specified by this document.
4. If time window enforcement is needed, then perform the this step and the next two steps. Append the string "&ET=". Note in the case of re-signing an URI, the attribute is carried over from the received Signed URI.
5. Get the current time in seconds since epoch (as an integer). Add the validity time in seconds as an integer. Note in the case of re-signing an URI, the value MUST remain the same as the received Signed URI.
6. Convert this integer to a string and append to the message.
7. If client IP address enforcement is needed, then perform this step and the next step. Append the string "&CIP=". Note in the case of re-signing an URI, the attribute is carried over from the received Signed URI.
8. Convert the client's IP address in dotted decimal notation format (i.e. for IPv4 address) or canonical text representation

(for IPv6 address [RFC5952]) to a string and append to the message. Note in the case of re-signing an URI, the value MUST remain the same as the received Signed URI.

9. Depending on the type of key used to sign the URI, compute the message digest or digital signature for symmetric key or asymmetric keys, respectively.
 - A. For symmetric key, HMAC is used.
 - a. Obtain the shared key to be used for signing the URI.
 - b. If the key identifier needs to be specified, then perform this step and the next step. Append the string "&KID=".
 - c. Append the key identifier (e.g. "example:keys:123") needed by the entity to locate the shared key for validating the URI signature.
 - d. If the hash function for HMAC needs to be specified, then perform this step and the next step. Append the string "&HF=".
 - e. Append the string for the type of hash function (e.g. "MD5", "SHA-1"). Note that re-signing an URI MUST use the same hash function as the received Signed URI or one of the allowable hash functions designated by the CDNI metadata.
 - f. Append the string "&MD=".
 - g. The message contains the complete section of the URI that is protected. (e.g. "://example.com/content.mov?VER=1&ET=1209422976&CIP=10.0.0.1&KID=example:keys:123&HF=SHA-1&MD=").
 - h. Compute the message digest (note: this is the URI signature) using the HMAC algorithm with the shared key and message as the two inputs to the hash function which is specified by the "HF" attribute. Note: Computation happens before URL encoding with escape characters is performed for reserved characters.
 - i. Convert the message digest to its equivalent hexadecimal format.

- j. Append the string for the message digest (e.g. "://example.com/content.mov?VER=1&ET=1209422976&CIP=10.0.0.1&KID=example:keys:123&HF=SHA-1&MD=4fb1cladfl588fbellcc6a04c6e69f35").
- B. For asymmetric keys, EC DSA is used.
- a. Generate the EC private and public key pair. Store the EC public key in a location that's reachable for any entity that needs to validate the URI signature.
 - b. If the key identifier needs to be specified, then perform this step and the next step. Append the string "&KID=".
 - c. Append the key identifier (e.g. "http://example.com/public/keys/123") needed by the entity to locate the shared key for validating the URI signature. Note the Key ID URI contains only the "scheme name", "authority", and "path" parts.
 - d. Append the string "&DS=".
 - e. The message contains the complete section of the URI that is protected. (e.g. "://example.com/content.mov?VER=1&ET=1209422976&CIP=10.0.0.1&KID=http://example.com/public/keys/123&DS=").
 - f. Compute the message digest using SHA-1 (without a key) for the message.
 - g. Compute the digital signature (note: this is the URI signature) using the EC DSA algorithm with the private EC key and message digest (obtained in previous step) as inputs. Note: Computation happens before URL encoding with escape characters is performed for reserved characters.
 - h. Convert the digital signature to its equivalent hexadecimal format.
 - i. Append the string for the digital signature which contains the values for the 'r' and 's' parameters. The (r,s) pair is denoted by ':' (e.g. "://example.com/content.mov?VER=1&ET=1209422976&CIP=10.0.0.1&KID=http://example.com/public/keys/123&

```
DS=r:
CFB03EDB33810AB6C79EE3C47FBD86D227D702F25F66C01CF03F59F1
E005668D:s:
57ED0E8DF7E786C87E39177DD3398A7FB010E6A4C0DC8AA71331A929
A29EA24E" )
```

10. When packaging the URI Signing attributes is not necessary, generate the Signed URI by prepending the "scheme name" part to the message (e.g. `http://example.com/content.mov?VER=1&ET=1209422976&CIP=10.0.0.1&KID=http://example.com/public/keys/123&DS=r:CFB03EDB33810AB6C79EE3C47FBD86D227D702F25F66C01CF03F59F1E005668D:s:57ED0E8DF7E786C87E39177DD3398A7FB010E6A4C0DC8AA71331A929A29EA24E"`). Note: this is the Signed URI before URI Signing Package is applied.

Apply the URI Signing Package attribute by following the procedure below to generate the URL Signing token.

1. Generate the URI Signing token in this step and the next step. Remove the Original URI portion from the message to obtain all the URI Signing attributes, including the URI signature (e.g. "VER=1&ET=1209422976&CIP=10.0.0.1&KID=example:keys:123&HF=SHA-1&MD=4fb1cladfl1588fbellcc6a04c6e69f35").
2. Compute the URI Signing token using Base-64 Data Encoding [RFC4648] on the message (e.g. "VkVSPTEmRVQ9MTIwOTQyMjk3NiZSVA9MTAuMC4wLjEmS0lEPWZvb2JhcjprZXlzMjESeEtMSZNrd00ZmIxYzFhZGYxNTg4ZmJlMTFjYzZhMDRjNmU2OWYzNQ=="). Note: This is the value for the URI Signing token.
3. Append the URI Signing token to the Original URI in this step and the next three steps. Copy the entire Original URI into a buffer to hold the message.
4. Check if the URI already contains a query string. If not, append a "?" character. If yes, append an "&" character.
5. Append the string "URISigningToken=" to the message.
6. Append the URI Signing token to the message (e.g. "http://example.com/content.mov?URISigningToken=VkVSPTEmRVQ9MTIwOTQyMjk3NiZSVA9MTAuMC4wLjEmS0lEPWZvb2JhcjprZXlzMjESeEtMSZNrd00ZmIxYzFhZGYxNTg4ZmJlMTFjYzZhMDRjNmU2OWYzNQ="). Note: this is the completed

Signed URI.

4. Validating a URI Signature

The following steps are specified to validate a Signed URI. However, it is possible to use some other algorithm and implementation as long as the same result is achieved. Note that some steps are to be skipped if the corresponding URI Signing attribute is not embedded in the Signed URI. The absence of a given attribute indicates enforcement of its purpose is not necessary in the distribution policy.

1. Extract the value from "URISigningToken" attribute. This value is the encoded URI Signing token. If there are multiple instances of this attribute, the last one is used.
2. Decode the string using Base-64 Data Encoding [RFC4648] (or another encoding method specified by configuration or CDNI metadata) to obtain all the URI Signing attributes (e.g. "VER=1&ET=1209422976&CIP=10.0.0.1&KID=example:keys:123&HF=SHA-1&MD=4fb1cladfl588fbellcc6a04c6e69f35").
3. Extract the value from "VER" attribute if the attribute exists. Determine the version of the URI Signing algorithm used to process the Signed URI. If the attribute is not in the URI, then obtain the version number in another manner (e.g. configuration or CDNI metadata).
4. Extract the value from "CIP" attribute if the attribute exists. Validate that the request came from the same IP address as indicated in the "CIP" attribute. If the IP address is incorrect, then the request is denied.
5. Extract the value from "ET" attribute if the attribute exists. Validate that the request arrived before expiration time based on the "ET" attribute. If the time expired, then the request is denied.
6. Extract the value from "MD" attribute if the attribute exists. If there are multiple instances of this attribute, the last one is used. [Editor's note: remove sentence if URISigningToken is mandatory] The attribute indicates symmetric key is used.
7. Extract the value from "DS" attribute if the attribute exists. If there are multiple instances of this attribute, the last one is used. [Editor's note: remove sentence if URISigningToken is mandatory] The attribute indicates asymmetric key is used.

8. If neither "MD" or "DS" attribute is in the URI, then no URI signature exists and the request is denied.

Validate the URI signature for the Signed URI.

1. Copy the Original URI, excluding the "scheme name" part, into a buffer to hold the message for performing the operations below.
2. Remove the "URISigningToken" attribute from the message.
3. Append the decoded value from "URISigningToken" attribute (which contains all the URI Signing attributes).
4. Depending on the type of key used to sign the URI, validate the message digest or digital signature for symmetric key or asymmetric keys, respectively.
 - A. For symmetric key, HMAC algorithm is used.
 - a. Extract the value from the "KID" attribute if the attribute exists. Use the key identifier (e.g. "example:keys:123") to locate the shared key, which may be one of the keys available to use (i.e. set by configuration or CDNI metadata). If the attribute is not in the URI, then obtain the key in another manner (e.g. configuration or CDNI metadata).
 - b. Extract the value from the "HF" attribute if the attribute exists. Determine the type of hash function (e.g. "MD5", "SHA-1") to use for HMAC. If the attribute is not in the URI, then obtain the hash function type in another manner (e.g. configuration or CDNI metadata).
 - c. Extract the value from the "MD" attribute. This is the received message digest.
 - d. Convert the message digest to binary format. This will be used to compare with the computed value later.
 - e. Remove the value part of the "MD" attribute (but not the '=' character) from the message. The message is ready for validation of the message digest (e.g. "://example.com/content.mov?VER=1&ET=1209422976&CIP=10.0.0.1&KID=example:keys:123&HF=SHA-1&MD=").
 - f. Compute the message digest using the HMAC algorithm with the shared key and message as the two inputs to the hash

function which is specified by the "HF" attribute. Note: Computation happens after URL decoding of the escape characters is performed for reserved characters.

- g. Compare the result with the received message digest to validate the Signed URI.

B. For asymmetric keys, EC DSA is used.

- a. Extract the value from the "KID" attribute. Use the key identifier (e.g. "http://example.com/public/keys/123") to obtain the EC public key, which may be one of the keys available to use (i.e. set by configuration or CDNI metadata). If the attribute is not in the URI, then obtain the key in another manner (e.g. configuration or CDNI metadata).
- b. Extract the value from the "DS" attribute. This is the digital signature.
- c. Convert the digital signature to binary format. This will be used for verification later.
- d. Remove the value part of the "DS" attribute (but not the '=' character) from the message. The message is ready for validation of the digital signature (e.g. "://example.com/content.mov?VER=1&ET=1209422976&CIP=10.0.0.1&KID=http://example.com/public/keys/123&DS=").
- e. Compute the message digest using SHA-1 (without a key) for the message. Note: Computation happens after URL decoding of the escape characters is performed for reserved characters.
- f. Verify the digital signature using the EC DSA algorithm with the public EC key, received digital signature, and message digest (obtained in previous step) as inputs. This validates the Signed URI.

5. Considerations for CDNI Interfaces

Some of the CDNI Interfaces need enhancements to support URI Signing. As an example: A Downstream CDN that supports URI Signing needs to be able to advertise this capability to the Upstream CDN. The Upstream CDN needs to select a Downstream CDN based on such capability when the CSP requires access control to enforce its distribution policy

via URI Signing. Also, the Upstream CDN needs to be able to distribute via the CDNI Metadata interface the information necessary to allow the Downstream CDN to validate a Signed URI . Events that pertain to URI Signing (e.g. request denial or delivery after access authorization) need to be included in the logs communicated through the CDNI Logging interface (Editor's Note: Is this within the scope of the CDNI Logging Interface?).

5.1. CDNI Control Interface

URI Signing has no impact on this interface.

5.2. CDNI Footprint & Capabilities Advertisement Interface

The Downstream CDN advertises its capability to support URI Signing via the CDNI Footprint & Capabilities Advertisement Interface (FCI). The supported version of URI Signing needs to be included to allow for future extensibility.

TBD: To be taken into account by Footprint & Capabilities design team working on this area.

- o [Editor's note: Advertise the CDNI metadata for URI Signing capability?]

5.3. CDNI Request Routing Redirection Interface

[Editor's Note: Debate the approach of dCDN providing the Signed URI vs. uCDN performing the signing function. List the pros/cons of each approach for the CDNI Request Routing Redirection interface (RI). Offer recommendation?]

The two approaches:

1. Downstream CDN provides the Signed URI

- * Key distribution is not necessary
- * Downstream CDN can use any scheme for Signed URI as long as the security level meets the CSP's expectation

2. Upstream CDN signs the URI

- * Consistency with interactive request routing method
- * URI Signing works even when Downstream CDN does not have the signing function (which may be the case when the Downstream CDN operates only as a delivering CDN)

- * Upstream CDN can act as a conversion gateway for the requesting routing interface between Upstream CDN and CSP and request routing interface between Upstream CDN and Downstream CDN since these two interfaces may not be the same

TBD: CDNI Redirection Interface is work in progress.

5.4. CDNI Metadata Interface

The following CDNI Metadata objects are specified for URI Signing.

- o URI Signing enforcement flag. Specifically, this flag indicates if the access to content is subject to URI Signing. URI Signing requires the Downstream CDN to ensure that the URI must be signed and validated before content delivery. Otherwise, Downstream CDN does not perform validation regardless if URI is signed or not.
- o Designated key identifier used for URI Signing computation when the Signed URI does not contain the Key ID attribute
- o Allowable Key ID set that the Signed URI's Key ID attribute can reference
- o Designated hash function used for URI Signing computation when the Signed URI does not contain the Hash Function attribute
- o Allowable Hash Function set that the Signed URI's Hash Function attribute can reference
- o Allowable crypto algorithm set? [Editor's note: TBD]
- o Overwrite the default name for the URL Signing Attribute Set attribute
- o Designated version used for URI Signing computation when the Signed URI does not contain the VER attribute
- o Allowable version/algorithm set that the Signed URI's VER attribute can reference
- o Allowable set of Downstream CDNs that participate in URI Signing based on the symmetric key
- o Overwrite the default encoding method for URI Signing Attribute Set attribute

Note that the Key ID information is not needed if only one key is provided by the CSP or the Upstream CDN for the content item or set

of content items covered by the CDNI Metadata object. In the case of asymmetric keys, it's easy for any entity to sign the URI for a content with a private key and provide the public key in the Signed URI. This just confirms that the URI Signer authorized the delivery. But it's necessary for the URI Signer to be the content owner. So, the CDNI Metadata Interface MUST provide the public key for the content or information to authorize the received Key ID attribute.

TBD: CDNI Metadata Interface is work in progress.

5.5. CDNI Logging Interface

The Downstream CDN reports that enforcement of the access control was applied to the request for content delivery.

The following CDNI Logging field for URI Signing SHOULD be supported in the HTTP Request Logging Record as specified in CDNI Logging Interface [I-D.ietf-cdni-logging].

- o s-uri-signing:

- * format: 1DIGIT

- * field value: this characterises the uri signing validation performed by the Surrogate on the request. The

- + "0" : no uri signature validation performed

- + "1" : uri signature validation performed and validated

- + "2" : uri signature validation performed and rejected

- allowed values are:

- * occurrence: there MUST be zero or exactly one instance of this field.

[Editor's note: Need to log these URI signature validation events (e.g. invalid client IP address, expired signed URI, incorrect URI signature, successful validation)?]

TBD: CDNI Logging interface is work in progress.

6. Detailed URI Signing Operation

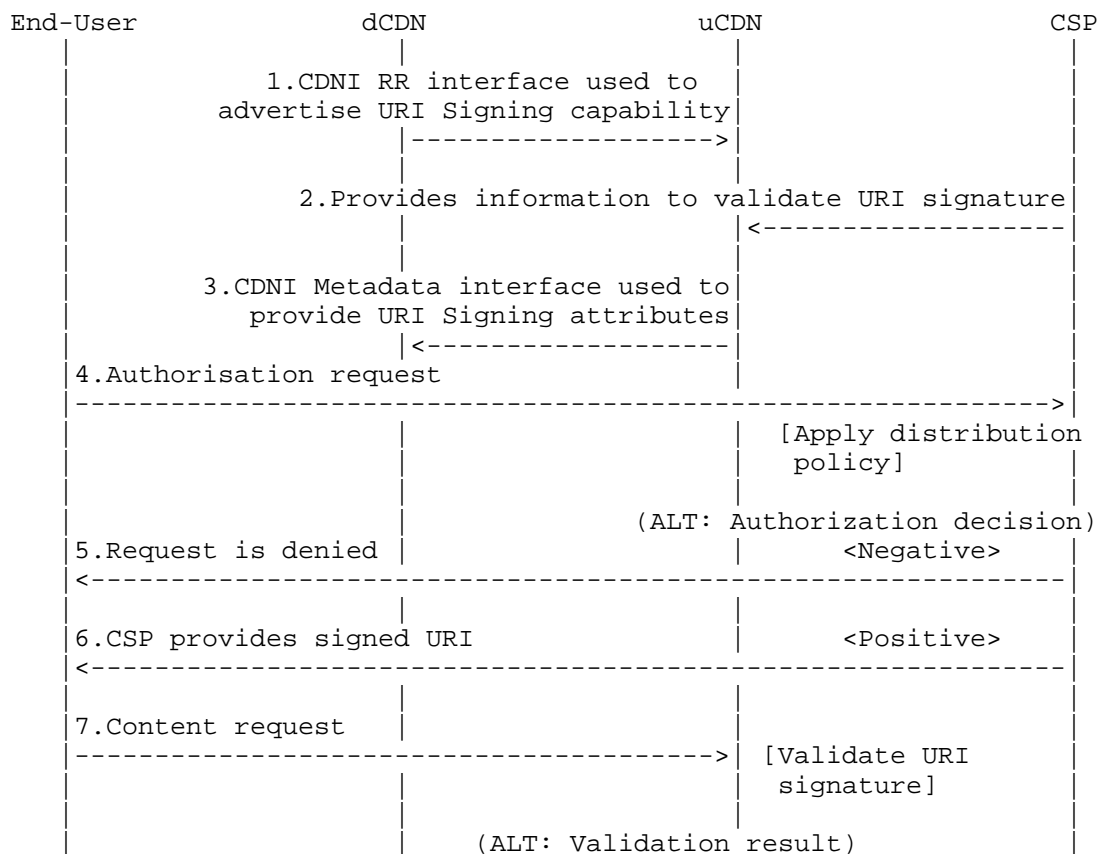
URI Signing supports both HTTP-based and DNS-based request routing. HMAC [RFC2104] defines a hash-based message authentication code

allowing two parties that share a symmetric key or asymmetric keys to establish the integrity and authenticity of a set of information (e.g. a message) through a cryptographic hash function.

6.1. HTTP Redirection

For HTTP-based request routing, HMAC is applied to a set of information that is unique to a given end user content request using key information that is specific to a pair of adjacent CDNI hops (e.g. between the CSP and the Authoritative CDN, between the Authoritative CDN and a Downstream CDN). This allows a CDNI hop to ascertain the authenticity of a given request received from a previous CDNI hop.

The URI signing scheme described below is based on the following steps (assuming HTTP redirection, iterative request routing and a CDN path with two CDNs). Note that Authoritative CDN and Upstream CDN are used exchangeably.



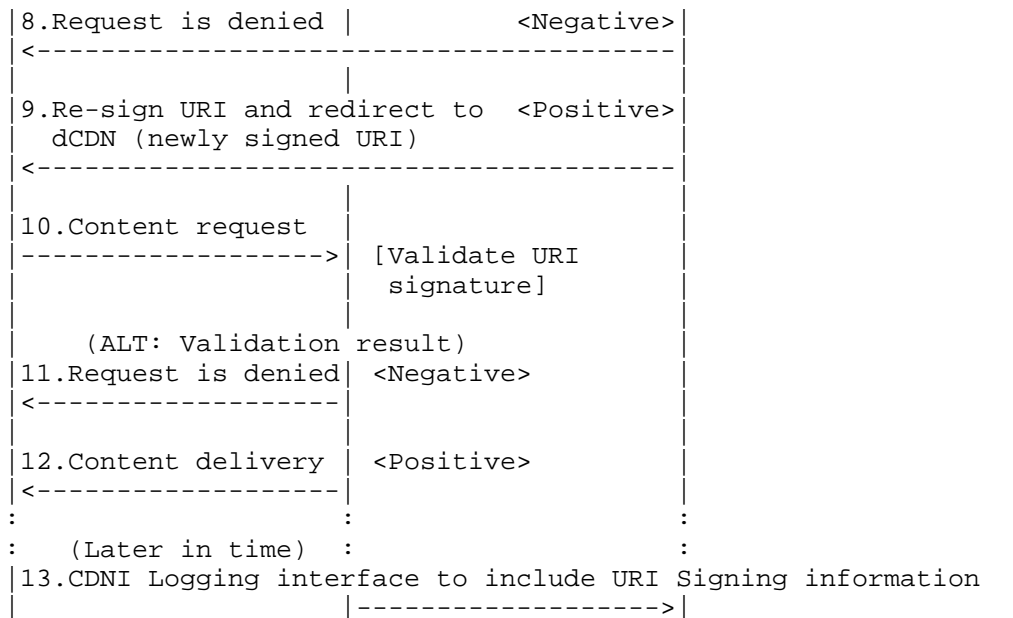


Figure 3: HTTP-based Request Routing with URI Signing

1. Using the CDNI Request Routing/Footprint & Capabilities Advertisement interface, the Downstream CDN advertises its capabilities including URI Signing support to the Authoritative CDN.
2. CSP provides to the Authoritative CDN the information needed to validate URI signatures from that CSP. For example, this information may include a hashing function, algorithm, and a key value.
3. Using the CDNI Metadata interface, the Authoritative CDN communicates to a Downstream CDN the information needed to validate URI signatures from the Authoritative CDN for the given CSP. For example, this information may include a hashing algorithm and private key corresponding to the trust relationship between the Authoritative CDN and the Downstream CDN.
4. When a UA requests a piece of protected content from the CSP, the CSP makes a specific authorization decision for this unique request based on its arbitrary distribution policy
5. If the authorization decision is negative, the CSP rejects the request.

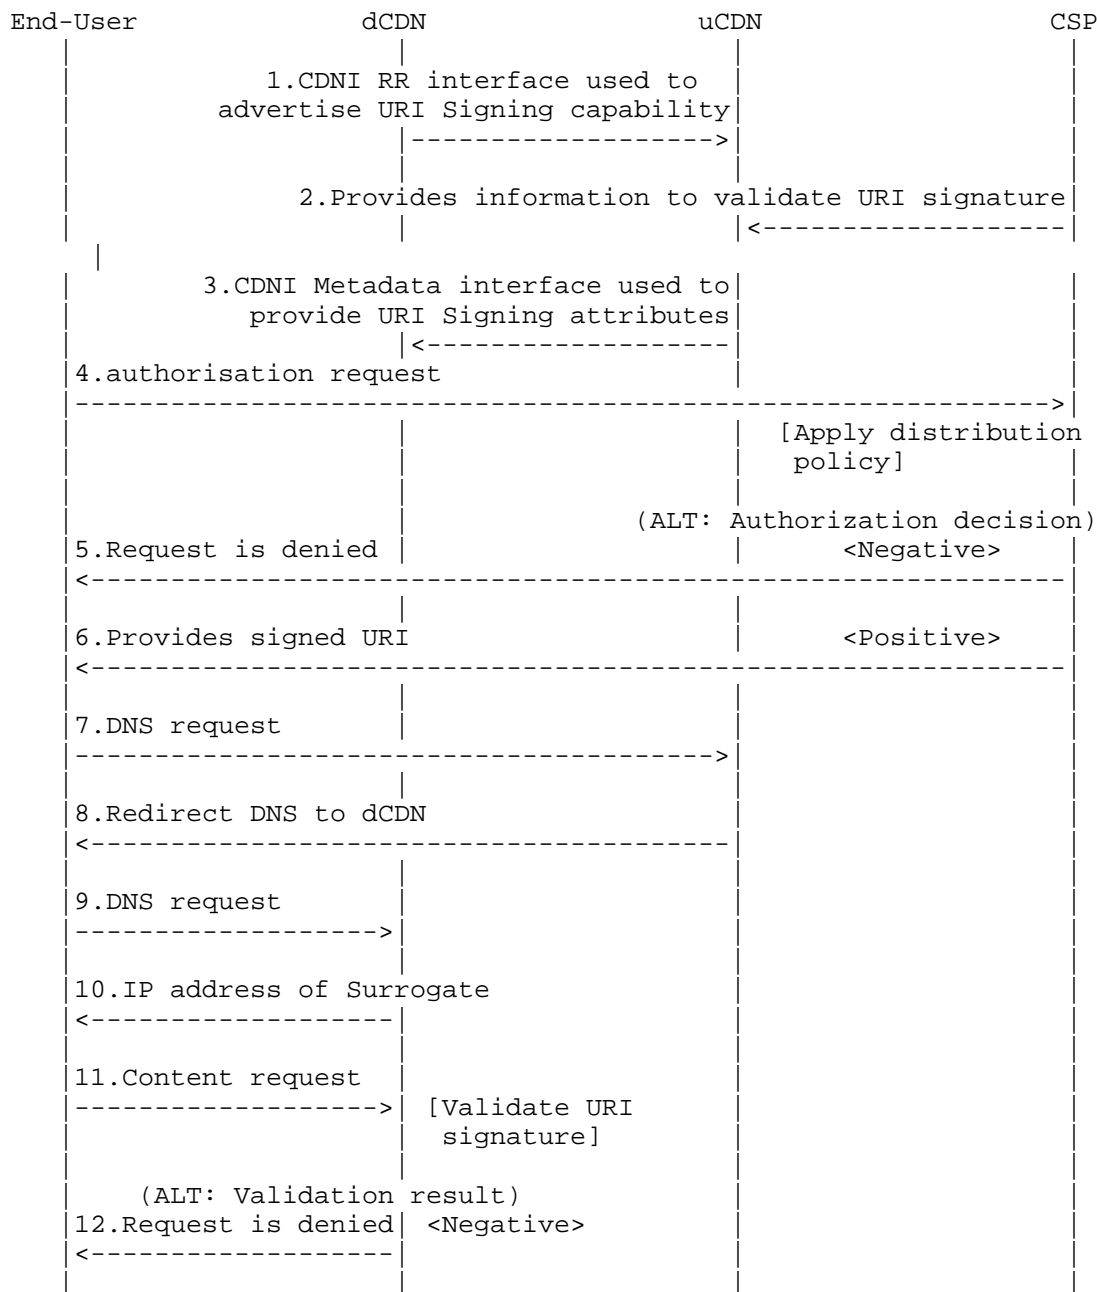
6. If the authorization decision is positive, the CSP computes a Signed URI that is based on unique parameters of that request and conveys it to the end user as the URI to use to request the content.
7. On receipt of the corresponding content request, the authoritative CDN validates the URI Signature in the URI using the information provided by the CSP.
8. If the validation is negative, the authoritative CDN rejects the request
9. If the validation is positive, the authoritative CDN computes a Signed URI that is based on unique parameters of that request and provides to the end user as the URI to use to further request the content from the Downstream CDN
10. On receipt of the corresponding content request, the Downstream CDN validates the URI Signature in the Signed URI using the information provided by the Authoritative CDN in the CDNI Metadata
11. If the validation is negative, the Downstream CDN rejects the request and sends an error code (e.g. 403) in the HTTP response.
12. If the validation is positive, the Downstream CDN serves the request and delivers the content.
13. At a later time, Downstream CDN reports logging events that includes URI signing information.

With HTTP-based request routing, URI Signing matches well the general chain of trust model of CDNI both with symmetric key and asymmetric keys because the key information only need to be specific to a pair of adjacent CDNI hops.

6.2. DNS Redirection

For DNS-based request routing, the CSP and Authoritative CDN must agree on a trust model appropriate to the security requirements of the CSP's particular content. Use of asymmetric public/private keys allows for unlimited distribution of the public key to downstream CDNs. However, if a shared secret key is preferred, then the CSP may want to restrict the distribution of the key to a (possibly empty) subset of trusted Downstream CDNs. Authorized Delivery CDNs need to obtain the key information to validate the Signed UR, which is computed by the CSP based on its distribution policy.

The URI signing scheme described below is based on the following steps (assuming iterative DNS request routing and a CDN path with two CDNs). Note that Authoritative CDN and Upstream CDN are used exchangeably.



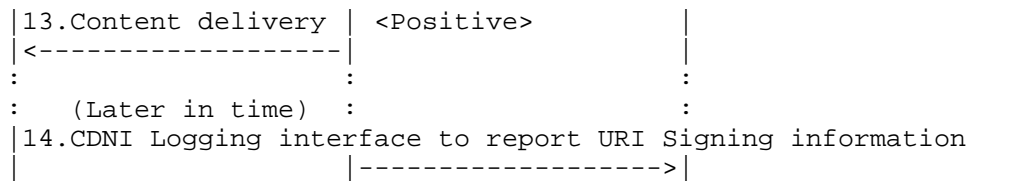


Figure 4: DNS-based Request Routing with URI Signing

1. Using the CDNI Request Routing interface, the Downstream CDN advertises its capabilities including URI Signing support to the Authoritative CDN.
2. CSP provides to the Authoritative CDN the information needed to validate cryptographic signatures from that CSP. For example, this information may include a hash function, algorithm, and a key.
3. Using the CDNI Metadata interface, the Authoritative CDN communicates to a Downstream CDN the information needed to validate cryptographic signatures from the CSP. In the case of symmetric key, the Authoritative CDN checks if the Downstream CDN is allowed by CSP to obtain the shared secret key.
4. When a UA requests a piece of protected content from the CSP, the CSP makes a specific authorization decision for this unique request based on its arbitrary distribution policy.
5. If the authorization decision is negative, the CSP rejects the request
6. If the authorization decision is positive, the CSP computes a cryptographic signature that is based on unique parameters of that request and includes it in the URI provided to the end user to request the content.
7. End user sends DNS request to the authoritative CDN.
8. On receipt of the DNS request, the authoritative CDN redirects the request to the Downstream CDN.
9. End user sends DNS request to the Downstream CDN.
10. On receipt of the DNS request, the Downstream CDN responds with IP address of one of its Surrogates.
11. On receipt of the corresponding content request, the Downstream CDN validates the cryptographic signature in the URI using the

information provided by the Authoritative CDN in the CDNI Metadata

12. If the validation is negative, the Downstream CDN rejects the request and sends an error code (e.g. 403) in the HTTP response.
13. If the validation is positive, the Downstream CDN serves the request and delivers the content.
14. At a later time, Downstream CDN reports logging events that includes URI signing information.

With DNS-based request routing, URI Signing matches well the general chain of trust model of CDNI when used with asymmetric keys because the only key information that need to be distributed across multiple CDNI hops including non-adjacent hops is the public key, that is generally not confidential.

With DNS-based request routing, URI Signing does match well the general chain of trust model of CDNI when used with symmetric keys because the symmetric key information needs to be distributed across multiple CDNI hops including non-adjacent hops. This raises a security concern for applicability of URI Signing with symmetric keys in case of DNS-based inter-CDN request routing.

7. HTTP Adaptive Streaming

The authors note that in order to perform URI signing for individual content segments of HTTP Adaptive Bitrate content, specific URI signing mechanisms are needed. Such mechanisms are currently out-of-scope of this document. More details on this topic is covered in Models for HTTP-Adaptive-Streaming-Aware CDNI [RFC6983].

8. IANA Considerations

[Editor note: (Is there a need to/How to) register official query string attribute keywords to be used for URI Signing? Need anything from IANA?]

This document requests IANA to create three new registries for the attributes (a.k.a. keywords) and their defined values in the URI Signing token.

This document highlights the use of the following query string attribute in the URI to support URI Signing. There is no intention to claim any query string attribute for URI beyond the CDNI URI

Signing context. That means the entities that sign the URI or validate the URI signature comply to the keyword specified in the query string for the URI Signing function only when URI Signing is used and only in the context of CDNI.

The following Enforcement Attributes names are allocated:

- o ET (Expiry time)
- o CIP (Client IP address)

The following Signature Computation Attributes names are allocated:

- o VER (Version): 1(Base)
- o KID (Key ID)
- o HF (Hash Function): "MD5", "SHA1"

The following URI Signature Attributes names are allocated:

- o MD (Message Digest)
- o DS (Digital Signature)

The following URI Signing Package Attribute names are allocated:

- o URI Signing Package (URI Signing token)

The IANA is requested to allocate a new entry to the CDNI Logging Field Names Registry as specified in CDNI Logging Interface [I-D.ietf-cdni-logging] in accordance to the "Specification Required" policy [RFC5226]

- o s-url-signing

9. Security Considerations

This document describes the concept of URI Signing and how it can be used to provide access authorization in the case of interconnected CDNs (CDNI). The primary goal of URI Signing is to make sure that only authorized UAs are able to access the content, with a Content Service Provider (CSP) being able to authorize every individual request. It should be noted that URI Signing is not a content protection scheme; if a CSP wants to protect the content itself, other mechanisms, such as DRM, are more appropriate.

In general it holds that the level of protection against illegitimate access can be increased by including more Enforcement Attributes in the URI. The current version of this document includes attributes for enforcing Client IP Address and Expiration Time, however this list can be extended with other, more complex, attributes that are able to provide some form of protection against some of the vulnerabilities highlighted below.

That said, there are a number of aspects that limit the level of security offered by URI signing and that anybody implementing URI signing should be aware of.

Replay attacks: Any (valid) Signed URI can be used to perform replay attacks. The vulnerability to replay attacks can be reduced by picking a relatively short window for the Expiration Time attribute, although this is limited by the fact that any HTTP-based request needs a window of at least a couple of seconds to prevent any sudden network issues from preventing legitimate UAs access to the content. One way to reduce exposure to replay attacks is to include in the URI a unique one-time access ID. Whenever the Downstream CDN receives a request with a given unique access ID, it adds that access ID to the list of 'used' IDs. In the case an illegitimate UA tries to use the same URI through a replay attack, the Downstream CDN can deny the request based on the already-used access ID.

Illegitimate client behind a NAT: In cases where there are multiple users behind the same NAT, all users will have the same IP address from the point of view of the Downstream CDN. This results in the Downstream CDN not being able to distinguish between the different users based on Client IP Address and illegitimate users being able to access the content. One way to reduce exposure to this kind of attack is to not only check for Client IP but also for other attributes that can be found in the HTTP headers.

TBD: ...

The shared key between CSP and Authoritative CDN may be distributed to Downstream CDNs - including cascaded CDNs. Since this key can be used to legitimately sign a URL for content access authorization, it's important to know the implications of a compromised shared key.

[Editor's note: Threat model cover in the Security section - Prevent client from spoofing URI (Ray) - Security implications - The scope of protection by URI Signing - Protects against DoS (network bandwidth and other nodes besides the edge cache); limits the time window.]

10. Privacy

The privacy protection concerns described in CDNI Logging Interface [I-D.ietf-cdni-logging] apply when the client's IP address (CIP attribute) is embedded in the Signed URI. In this case, all the CDNI attributes MUST be removed from the logging record when anonymization is enabled.

11. Acknowledgements

The authors would like to thank the following people for their contributions in reviewing this document and providing feedback: Kevin Ma, Ben Niven-Jenkins, Thierry Magnien, Dan York, Bhaskar Bhupalam, Matt Caulfield, and Samuel Rajakumar .

12. References

12.1. Normative References

- [I-D.ietf-cdni-logging]
Faucheur, F., Bertrand, G., Oprescu, I., and R. Peterkofsky, "CDNI Logging Interface", draft-ietf-cdni-logging-05 (work in progress), July 2013.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [RFC6707] Niven-Jenkins, B., Le Faucheur, F., and N. Bitar, "Content Distribution Network Interconnection (CDNI) Problem Statement", RFC 6707, September 2012.

12.2. Informative References

- [I-D.ietf-cdni-framework]
Peterson, L. and B. Davie, "Framework for CDN Interconnection", draft-ietf-cdni-framework-04 (work in progress), August 2013.
- [I-D.ietf-cdni-requirements]
Leung, K. and Y. Lee, "Content Distribution Network Interconnection (CDNI) Requirements", draft-ietf-cdni-requirements-09 (work in progress), July 2013.
- [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-10 (work in progress), August 2012.
- [RFC2104] Krawczyk, H., Bellare, M., and R. Canetti, "HMAC: Keyed-Hashing for Message Authentication", RFC 2104, February 1997.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005.
- [RFC4648] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", RFC 4648, October 2006.
- [RFC5952] Kawamura, S. and M. Kawashima, "A Recommendation for IPv6 Address Text Representation", RFC 5952, August 2010.
- [RFC6983] van Brandenburg, R., van Deventer, O., Le Faucheur, F., and K. Leung, "Models for HTTP-Adaptive-Streaming-Aware Content Distribution Network Interconnection (CDNI)", RFC 6983, July 2013.

Authors' Addresses

Kent Leung
Cisco Systems
3625 Cisco Way
San Jose 95134
USA

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

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

Bill Downey
Verizon Labs
60 Sylvan Road
Waltham, Massachusetts 02451
USA

Phone: +1 781 466 2475
Email: william.s.downey@verizon.com

Ray van Brandenburg
TNO
Brassersplein 2
Delft, 2612CT
the Netherlands

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

Scott Leibrand
Limelight Networks
222 S Mill Ave
Tempe, AZ 85281
USA

Phone: +1 360 419 5185
Email: sleibrand@llnw.com

CDNI
Internet-Draft
Intended status: Standards Track
Expires: September 5, 2014

K. Leung
F. Le Faucheur
Cisco Systems
B. Downey
Verizon Labs
R. van Brandenburg
TNO
S. Leibrand
Limelight Networks
March 4, 2014

URI Signing for CDN Interconnection (CDNI)
draft-leung-cdni-uri-signing-05

Abstract

This document describes how the concept of URI signing supports the content access control requirements of CDNI and proposes a URI signing scheme.

The proposed URI signing method specifies the information needed to be included in the URI and the algorithm used to authorize and to validate access requests for the content referenced by the URI. Some of the information may be accessed by the CDN via configuration or CDNI metadata.

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 September 5, 2014.

Copyright Notice

Copyright (c) 2014 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	3
1.2. Background on URI Signing	4
1.3. CDNI URI Signing Overview	6
1.4. URI Signing in a non-CDNI context	8
2. Signed URI Information Elements	8
2.1. Enforcement Information Elements	10
2.2. Signature Computation Information Elements	11
2.3. URI Signature Information Elements	12
2.4. URI Signing Package Attribute	13
3. Creating the Signed URI	14
3.1. Calculating the URI Signature	14
3.2. Packaging the URI Signature	17
4. Validating a URI Signature	18
4.1. Information element validation	18
4.2. Signature validation	19
5. Relationship with CDNI Interfaces	21
5.1. CDNI Control Interface	22
5.2. CDNI Footprint & Capabilities Advertisement Interface	22
5.3. CDNI Request Routing Redirection Interface	22
5.4. CDNI Metadata Interface	23
5.5. CDNI Logging Interface	24
6. URI Signing Message Flow	24
6.1. HTTP Redirection	25
6.2. DNS Redirection	27
7. HTTP Adaptive Streaming	30
8. IANA Considerations	30
9. Security Considerations	31
10. Privacy	33
11. Acknowledgements	33
12. References	33
12.1. Normative References	33
12.2. Informative References	33
Authors' Addresses	34

1. Introduction

This document describes the concept of URI Signing and how it can be used to provide access authorization in the case of interconnected CDNs (CDNI). The primary goal of URI Signing is to make sure that only authorized User Agents (UAs) are able to access the content, with a Content Service Provider (CSP) being able to authorize every individual request. It should be noted that URI Signing is not a content protection scheme; if a CSP wants to protect the content itself, other mechanisms, such as DRM, are more appropriate.

The overall problem space for CDN Interconnection (CDNI) is described in CDNI Problem Statement [RFC6707]. In this document, along with the CDNI Requirements [I-D.ietf-cdni-requirements] document and the CDNI Framework [I-D.ietf-cdni-framework] the need for interconnected CDNs to be able to implement an access control mechanism that enforces the CSP's distribution policy is described.

Specifically, CDNI Framework [I-D.ietf-cdni-framework] states:

"The CSP may also trust the CDN operator to perform actions such as ..., and to enforce per-request authorization performed by the CSP using techniques such as URI signing."

In particular, the following requirement is listed in CDNI Requirements [I-D.ietf-cdni-requirements]:

"MI-16 [HIGH] The CDNI Metadata Distribution interface shall allow signaling of authorization checks and validation that are to be performed by the surrogate before delivery. For example, this could potentially include:

- * need to validate URI signed information (e.g. Expiry time, Client IP address)."

This document proposes a URI Signing scheme that allows Surrogates in interconnected CDNs to enforce a per-request authorization performed by the CSP. Splitting the role of performing per-request authorization by CSP and the role of validation of this authorization by the CDN allows any arbitrary distribution policy to be enforced across CDNs without the need of CDNs to have any awareness of the actual CSP distribution policy.

1.1. Terminology

This document uses the terminology defined in CDNI Problem Statement [RFC6707].

This document also uses the terminology of Keyed-Hashing for Message Authentication (HMAC) [RFC2104] including the following terms (reproduced here for convenience):

- o MAC: message authentication code.
- o HMAC: Hash-based message authentication code (HMAC) is a specific construction for calculating a MAC involving a cryptographic hash function in combination with a secret key.
- o HMAC-SHA1: HMAC instantiation using SHA-1 as the cryptographic hash function.
- o HMAC-MD5: HMAC instantiation using MD5 as the cryptographic hash function.

In addition, the following terms are used throughout this document:

- o URI Signature: Message digest or digital signature that is computed with an algorithm for protecting the URI.
- o Original URI: The URI before URI Signing is applied.
- o Signed URI: Any URI that contains a URI Signature.
- o Target CDN URI: Embedded URI created by the CSP to direct UA towards the Upstream CDN. The Target CDN URI can be signed by the CSP and verified by the Upstream CDN.
- o Redirection URI: URI created by the Upstream CDN to redirect UA towards the Downstream CDN. The Redirection URI can be signed by the Upstream CDN and verified by the Downstream CDN. In a cascaded CDNI scenario, there can be more than one Redirection URI.

1.2. Background on URI Signing

The next section provides an overview of how URI Signing works in a CDNI environment. As background information, URI Signing is first explained in terms of a single CDN delivering content on behalf of a CSP.

A CSP and CDN are assumed to have a trust relationship that enables the CSP to authorize access to a content item by including a set of attributes in the URI before redirecting a UA to the CDN. Using these attributes, it is possible for a CDN to check an incoming content request to see whether it was authorized by the CSP (e.g. based on the UA's IP address or a time window). Of course, the

attributes need to be added to the URI in a way that prevents a UA from changing the attributes, thereby leaving the CDN to think that the request was authorized by the CSP when in fact it wasn't. For this reason, a URI Signing mechanism includes in the URI a message digest or digital signature that allows a CDN to check the authenticity of the URI. The message digest or digital signature can be calculated based on a shared secret between the CSP and CDN or using CSP's asymmetric public/private key pair, respectively.

Figure 1, shown below, presents an overview of the URI Signing mechanism in the case of a CSP with a single CDN. When the UA browses for content on CSP's website (#1), it receives HTML web pages with embedded content URIs. Upon requesting these URIs, the CSP redirects to a CDN, creating a Target CDN URI (#2) (alternatively, the Target CDN URI itself is embedded in the HTML). The Target CDN URI is the Signed URI which may include the IP address of the UA and/or a time window and always contains the URI Signature which is generated by the CSP using the shared secret or a private key. Once the UA receives the response with the embedded URI, it sends a new HTTP request using the embedded URI to the CDN (#3). Upon receiving the request, the CDN checks to see if the Signed URI is authentic by verifying the URI signature. In addition, it checks whether the IP address of the HTTP request matches that in the Signed URI and if the time window is still valid. After these values are confirmed to be valid, the CDN delivers the content (#4).

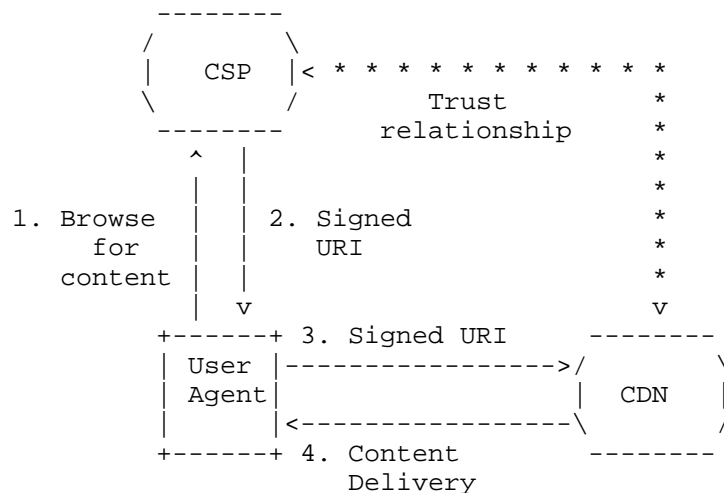


Figure 1: Figure 1: URI Signing in a CDN Environment

1.3. CDNI URI Signing Overview

In a CDNI environment, URI Signing operates the same way in the initial steps #1 and #2 but the later steps involve multiple CDNs in the process of delivering the content. The main difference from the single CDN case is a redirection step between the Upstream CDN and the Downstream CDN. In step #3, UA may send HTTP request or DNS request. Depending on whether HTTP-based or DNS-based request routing is used, the Upstream CDN responds by directing the UA towards the Downstream CDN using either a Redirection URI (which is a Signed URI generated by the Upstream CDN) or a DNS reply, respectively (#4). Once the UA receives the response, it sends the Redirection URI/Target CDN URI to the Downstream CDN (#5). The received URI is validated by the Downstream CDN before delivering the content (#6). This is depicted in the figure below. Note: The CDNI call flows are covered in Detailed URI Signing Operation (Section 6).

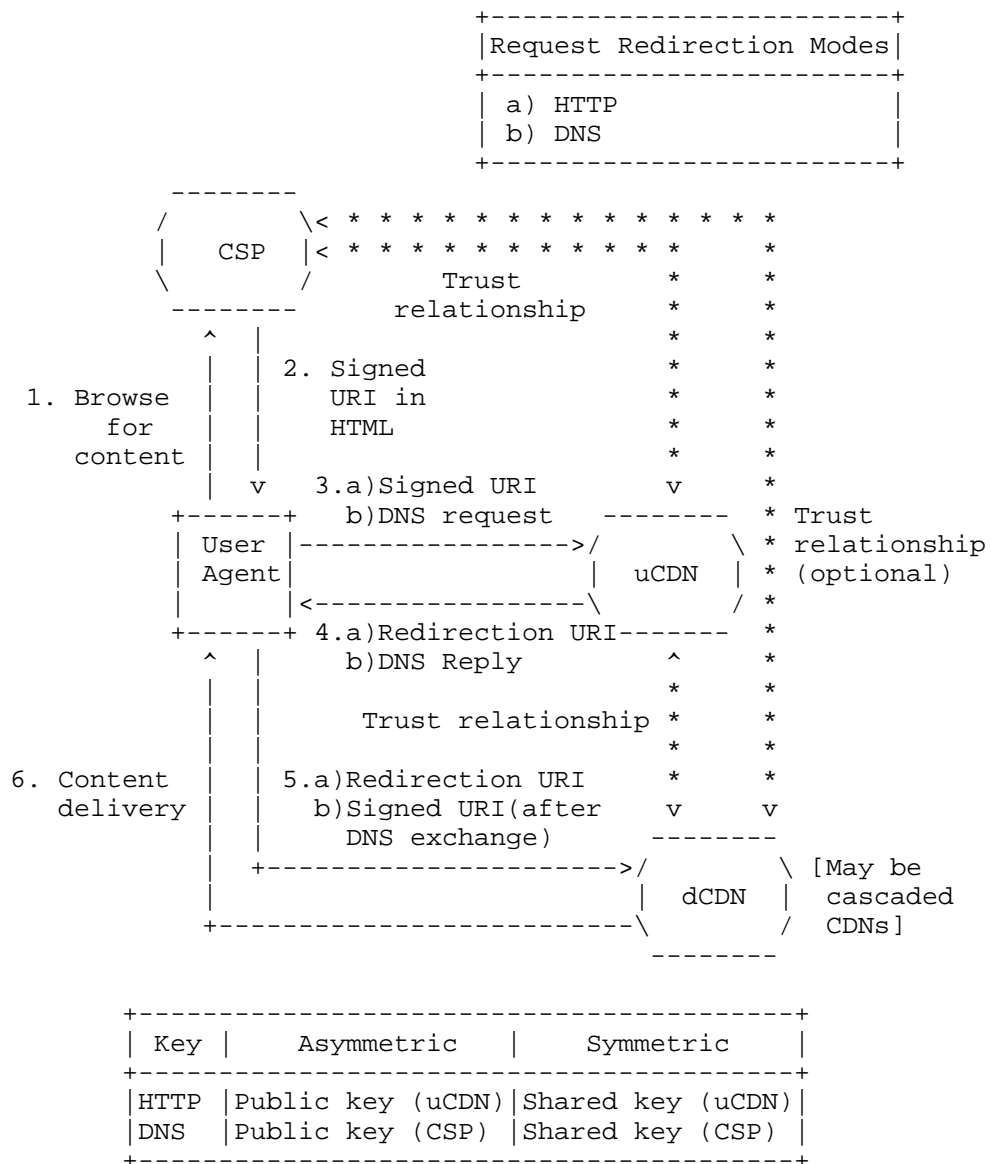


Figure 2: URI Signing in a CDNI Environment

The trust relationships between CSP, Upstream CDN, and Downstream CDN have direct implications for URI Signing. In the case shown in Figure 2, the CDN that the CSP has a trust relationship with is the Upstream CDN. The delivery of the content may be delegated to the Downstream CDN, which has a relationship with the Upstream CDN but

may have no relationship with the CSP.

In CDNI, there are two methods for request routing: DNS-based and HTTP-based. For DNS-based request routing, the Signed URI (i.e. Target CDN URI) provided by the CSP reaches the Downstream CDN directly. In the case where the Downstream CDN does not have a trust relationship with the CSP, this means that only an asymmetric public/private key method can be used for computing the URI Signature because the CSP and Downstream CDN are not able to exchange symmetric shared secret keys. Since the CSP is unlikely to have relationships with all the Downstream CDNs that are delegated to by the Upstream CDN, the CSP may choose to allow the Authoritative CDN to redistribute the shared key to a subset of their Downstream CDNs .

For HTTP-based request routing, the Signed URI (i.e. Target CDN URI) provided by the CSP reaches the Upstream CDN. After this URI has been verified to be correct by the Upstream CDN, the Upstream CDN creates and signs a new Redirection URI to redirect the UA to the Downstream CDN. Since this new URI also has a new URI Signature, this new signature can be based around the trust relationship between the Upstream CDN and Downstream CDN, and the relationship between the Downstream CDN and CSP is not relevant. Given the fact that such a relationship between Upstream CDN and Downstream CDN always exists, both asymmetric public/private keys and symmetric shared secret keys can be used for URI Signing. Note that the signed Redirection URI SHOULD maintain the same level of security as the original Signed URI.

1.4. URI Signing in a non-CDNI context

While the URI signing scheme defined in this document was primarily created for the purpose of allowing URI Signing in CDNI scenarios, e.g. between a uCDN and a dCDN or between a CSP and a dCDN, there is nothing in the defined URI Signing scheme that precludes it from being used in a non-CDNI context. As such, the described mechanism could be used in a single-CDN scenario such as shown in Figure 1 in Section 1.2, for example to allow a CSP that uses different CDNs to only have to implement a single URI Signing mechanism.

2. Signed URI Information Elements

The concept behind URI Signing is based on embedding in the Target CDN URI/Redirection URI a number of information elements that can be validated to ensure the UA has legitimate access to the content. These information elements are appended, in an encapsulated form, to the original URI.

For the purposes of the URI signing mechanism described in this document, three types of information elements may be embedded in the URI:

- o Enforcement Information Elements: Information Elements that are used to enforce a distribution policy defined by the CSP. Examples of enforcement attributes are IP address of the UA and time window.
- o Signature Computation Information Elements: Information Elements that are used by the CDN to verify the URI signature embedded in the received URI. In order to verify a URI Signature, the CDN requires some information elements that describe how the URI Signature was generated. Examples of Signature Computation Elements include the used HMACs hash function and/or the key identifier.
- o URI Signature Information Elements: The information elements that carry the actual message digest or digital signature representing the URI signature used for checking the integrity and authenticity of the URI. A typical Signed URI will only contain one embedded URI Signature Information Element.

In addition, the this document specifies the following URI attribute:

- o URI Signing Package Attribute: The URI attribute that encapsulates all the URI Signing information elements in an encoded format. Only this attribute is exposed in the Signed URI as a URI query parameter.

If the UA or another entity needs to add one or more attributes to the Signed URI for purposes other than URI Signing, those attributes MUST be appended after the URI Signing Packacke Attribute. Any attributes appended in such way after the URI Signature has been calculated are not validated for the purpose of content access authorization. Note that adding any such attributes to the Signed URI before the URI Signing Packacke Attribute will cause the URI Signing validation to fail.

Two types of keys can be used for URI Signing: asymmetric keys and symmetric keys. Asymmetric keys are based on a public/private key pair mechanism and always contain a private key only known to the entity signing the URI (either CSP or uCDN) and a public key for the verification of the Signed URI. With symmetric keys, the same key is used by both the signing entity for signing the URI as well as by the validating entity for validating the Signed URI. Regardless of the type of keys used, the validating entity has to obtain the key (either the public or the symmetric key). There are very different

requirements for key distribution (out of scope of this document) with asymmetric keys and with symmetric keys. Key distribution for symmetric keys requires confidentiality to prevent another party from getting access to the key, since it could then generate valid Signed URIs for unauthorized requests. Key distribution for asymmetric keys does not require confidentiality since public keys can typically be distributed openly (because they cannot be used for URI signing) and private keys are kept by the URI signing function.

2.1. Enforcement Information Elements

This section identifies the set of information elements that may be needed to enforce the CSP distribution policy. New information elements may be introduced in the future to extend the capabilities of the distribution policy.

In order to provide flexibility in distribution policies to be enforced, the exact subset of information elements used in the URI Signature of a given request is a deployment decision. The defined keyword for each information element is specified in parenthesis below.

The following information elements are used to enforce the distribution policy:

- o Expiry Time (ET) [optional] - Time when the Signed URI expires. This is represented as an integer denoting the number of seconds since midnight 1/1/1970 UTC (i.e. UNIX epoch). The request is rejected if the received time is later than this timestamp. Note: The time, including time zone, on the entities that generate and validate the signed URI need to be in sync (e.g. NTP is used).
- o Client IP (CIP) [optional] - IP address of the client for which this Signed URI is generated. This is represented in dotted decimal format for IPv4 or canonical text representation for IPv6 address [RFC5952] . The request is rejected if sourced from a client with a different IP address.

The Expiry Time Information Element ensures that the content authorization expires after a predetermined time. This limits the time window for content access and prevents replay of the request beyond the authorized time window.

The Client IP Information Element is used to restrict content access to a particular User Agent, based on its IP address for whom the content access was authorized.

Note: See the Security Considerations (Section 9) section on the

limitations of using an expiration time and client IP address for distribution policy enforcement.

2.2. Signature Computation Information Elements

This section identifies the set of information elements that may be needed to verify the URI (signature). New information elements may be introduced in the future if new URI signing algorithms are developed.

The defined keyword for each information element is specified in parenthesis below.

The following information elements are used to validate the URI by recreating the URI Signature.

- o Version (VER) [optional] - An integer used for identifying the version of URI signing method. If this Information Element is not present in the URI Signing Package Attribute, the default version is 1.
- o Key ID (KID) [optional] - A string used for obtaining the key (e.g. database lookup, URI reference) which is needed to validate the URI signature.
- o Hash Function (HF) [optional] - A string used for identifying the hash function to compute the URI signature (e.g. "MD5", "SHA-1", "SHA-256", "SHA-3") with HMAC. If this Information Element is not present in the URI Signing Package Attribute, the default hash function is SHA-1.
- o Digital Signature Algorithm (DSA) [optional] - Algorithm used to calculate the Digital Signature (e.g. "RSA", "DSA", "EC-DSA"). If this Information Element is not present in the URI Signing Package Attribute, the default is EC-DSA.

The Version Information Element indicates which version of URI signing scheme is used (including which attributes and algorithms are supported). The present document specifies Version 1. If the Version attribute is not present in the Signed URI, then the version is obtained from the CDNI metadata, else it is considered to have been set to the default value. More versions may be defined in the future.

The Key ID Information Element is used to retrieve the key which is needed as input to the algorithm for validating the Signed URI. The method used for obtaining the actual key from the reference included in the Key ID Information Element is outside the scope of this

document.

The Hash Function Information Element indicates the hash function to be used for HMAC-based message digest computation. The Hash Function Information Element is used in combination with the Message Digest Information Element defined in section Section 2.3.

The Digital Signature Algorithm Information Element indicates the digital signature function to be in the case asymmetric keys are used. The Digital Signature Algorithm Information Element is used in combination with the Digital Signature Information Element defined in section Section 2.3.

2.3. URI Signature Information Elements

This section identifies the set of information elements that carry the URI Signature that is used for checking the integrity and authenticity of the URI.

The defined keyword for each information element is specified in parenthesis below.

The following information elements are used to carry the actual URI Signature.

- o Message Digest (MD) [mandatory for symmetric key] - A string used for the message digest generated by the URI signing entity.
- o Digital Signature (DS) [mandatory for asymmetric keys] - A string used for the digital signature provided by the URI signing entity.

The Message Digest attribute contains the message digest used to validate the Signed URI when symmetric keys are used.

The Digital Signature attribute contains the digital signature used to verify the Signed URI when asymmetric keys are used.

In the case of symmetric key, HMAC algorithm is used for the following reasons: 1) Ability to use hash functions (i.e. no changes needed) with well understood cryptographic properties that perform well and for which code is freely and widely available, 2) Easy to replace the embedded hash function in case faster or more secure hash functions are found or required, 3) Original performance of the hash function is maintained without incurring a significant degradation, and 4) Simple way to use and handle keys.

In the case of asymmetric keys, Elliptic Curve Digital Signature Algorithm (EC DSA) - a variant of DSA - is used because of the

following reasons: 1) Key size is small while still offering good security, 2) Key is easy to store, and 3) Computation is faster than DSA or RSA.

2.4. URI Signing Package Attribute

The URI Signing Package Attribute is an encapsulation container for the URI Signing Information Elements defined in the previous sections. The URI Signing Information Elements are encoded and stored in this attribute. URI Signing Package Attribute is appended to the Original URI to create the Signed URI.

The primary advantage of the URI Signing Package Attribute is that it avoids having to expose the URI Signing Information Elements directly in the query string of the URI, thereby reducing the potential for a namespace collision space within the URI query string. A side benefit of the attribute is the obfuscation performed by the URI Signing Package Attribute hides the information (e.g. client IP address) from view of the common user, who is not aware of the encoding scheme. Obviously, this is not a security method since anyone who knows the encoding scheme is able to obtain the clear text. Note that any parameters appended to the query string after the URI Signing Package Attribute are not validated and hence do not affect URI Signing.

The following attribute is used to carry the encoded set of URI Signing attributes in the Signed URI.

- o URI Signing Package (URISigningPackage) - The encoded attribute containing all the CDNI URI Signing Information Elements used for URI Signing.

The URI Signing Package Attribute contains the URI Signing Information Elements in the Base-64 encoding with URL and Filename Safe Alphabet (a.k.a. "base64url") as specified in the Base-64 Data Encoding [RFC4648] document. The URI Signing Package Attribute is the only URI Signing attribute exposed in the Signed URI. The attribute MUST be the last parameter in the query string of the URI when the Signed URI is generated. However, a client or CDN may append other query parameters unrelated to URI Signing to the Signed URI. Such additional query parameters SHOULD NOT use the same name as the URI Signing Package Attribute to avoid namespace collision and potential failure of the URI Signing validation.

The parameter name of the URI Signing Package Attribute shall be defined in the CDNI Metadata interface. If the CDNI Metadata interface does not include a parameter name for the URI Signing Package Attribute, the parameter name is set by configuration ((out

of scope of this document).

3. Creating the Signed URI

The following procedure for signing a URI defines the algorithms in this version of URI Signing. Note that some steps may be skipped if the CSP does not enforce a distribution policy and the Enforcement Information Elements are therefore not necessary. A URI (as defined in URI Generic Syntax [RFC3986]) contains the following parts: scheme name, authority, path, query, and fragment. The entire URI except the "scheme name" part is protected by the URI signature. This allows the URI signature to be validated correctly in the case when a client performs a fallback to another scheme (e.g. HTTP) for a content item referenced by a URI with a specific scheme (e.g. RTSP). The benefit is that the content access is protected regardless of the type of transport used for delivery. If the CSP wants to ensure a specific protocol is used for content delivery, that information is passed by CDNI metadata. Note: Support for changing of the URL scheme requires that the default port is used, or that the protocols must both run on the same non-standard port.

The process of generating a Signed URI can be divided into two sets of steps: calculating the URI Signature and packaging the URI Signature and appending it to the Original URI. Note it is possible to use some other algorithm and implementation as long as the same result is achieved. An example for the Original URI, "http://example.com/content.mov", is used to clarify the steps.

3.1. Calculating the URI Signature

Calculate the URI Signature by following the procedure below.

1. Copy the Original URI, excluding the "scheme name" part, into a buffer to hold the message for performing the operations below.
2. Check if the URI already contains a query string. If not, append a "?" character. If yes, append an "&" character.
3. If the version is the default value, skip this step. Append the string "VER=". Append the string for the version number.
4. If time window enforcement is not needed, step 4 can be skipped.
 - A. If an attribute was added to the URI, append an "&" character. Append the string "ET=". Note in the case of re-signing a URI, the attribute is carried over from the received Signed URI.

- B. Get the current time in seconds since epoch (as an integer). Add the validity time in seconds as an integer. Note in the case of re-signing a URI, the value MUST remain the same as the received Signed URI.
 - C. Convert this integer to a string and append to the message.
5. If client IP enforcement is not needed, step 5 can be skipped.
- A. If an attribute was added to the URI, append an "&" character. Append the string "CIP=". Note in the case of re-signing a URI, the attribute is carried over from the received Signed URI.
 - B. Convert the client's IP address in dotted decimal notation format (i.e. for IPv4 address) or canonical text representation (for IPv6 address [RFC5952]) to a string and append to the message. Note in the case of re-signing an URI, the value MUST remain the same as the received Signed URI.
6. Depending on the type of key used to sign the URI, compute the message digest or digital signature for symmetric key or asymmetric keys, respectively.
- A. For symmetric key, HMAC is used.
 - 1. Obtain the shared key to be used for signing the URI.
 - 2. If the key identifier is provided by the CDNI metadata, skip this step. If an attribute was added to the URI, append an "&" character. Append the string "KID=". Append the key identifier (e.g. "example:keys:123") needed by the entity to locate the shared key for validating the URI signature.
 - 3. If the hash function for the HMAC uses the default value (SHA-1), skip this step. If an attribute was added to the URI, append an "&" character. Append the string "HF=". Append the string for the type of hash function. Note that re-signing a URI MUST use the same hash function as the received Signed URI or one of the allowable hash functions designated by the CDNI metadata.
 - 4. If an attribute was added to the URI, append an "&" character. Append the string "MD=". The message now contains the complete section of the URI that is protected (e.g. "://example.com/

content.mov?ET=1209422976&CIP=10.0.0.1&
KID=example:keys:123&MD=").

5. Compute the message digest using the HMAC algorithm with the shared key (e.g. "secretkey" and message as the two inputs to the hash function which is specified by the "HF" attribute.
6. Convert the message digest to its equivalent hexadecimal format.
7. Append the string for the message digest (e.g. "://example.com/
content.mov?ET=1209422976&CIP=10.0.0.1&
KID=example:keys:123&
MD=da58513e8b309c1e8a9695baceba629d180b50b8").

B. For asymmetric keys, EC DSA is used.

1. Generate the EC private and public key pair (e.g. private key is "8b5b417336492707a83836b02ceee55b3847be5ec1521e4949977b224950e708", public key is "04840b1be11cfd1404c2fc588d30150a4103cadcc4172e786bcafl5d7feeb6d246f7d8a91fa055cb10efb2f52860d1dlb2f339244e9ad79a23e10ed9b720f6157f"). Store the EC public key in a location that's reachable for any entity that needs to validate the URI signature.
2. If the key identifier is provided by the CDNI metadata, skip this step. If an attribute was added to the URI, append an "&" character. Append the string "KID=". Append the key identifier (e.g. "http://example.com/public/keys/123") needed by the entity to locate the shared key for validating the URI signature. Note the Key ID URI contains only the "scheme name", "authority", and "path" parts (i.e. query string is not allowed).
3. If the digital signature algorithm uses the default value (EC-DSA), skip this step. If an attribute was added to the URI, append an "&" character. Append the string "DSA=". Append the string denoting the digital signature function used.
4. If an attribute was added to the URI, append an "&" character. Append the string "DS=". The message now contains the complete section of the URI that is protected. (e.g. "://example.com/
content.mov?ET=1209422976&CIP=10.0.0.1&KID=http://

example.com/public/keys/123&DS=").

5. Compute the message digest using SHA-1 (without a key) for the message (e.g. message digest is "b95cb62f1d30ad03969619e9574a925fbfe9aeaf"). Note: The reason the digital signature calculated in the next step is calculated over the SHA-1 message digest, instead of over the cleartype message, is to reduce the length of the digital signature, and thereby the length of the URI Signing Package Attribute and the resulting Signed URI.
6. Compute the digital signature, using the EC-DSA algorithm by default or another algorithm if specified by the DSA Information Element, with the private EC key and message digest obtained in previous step as inputs.
7. Convert the digital signature to its equivalent hexadecimal format.
8. Append the string for the digital signature. In the case where EC-DSA algorithm is used, this string contains the values for the 'r' and 's' parameters, delimited by ':' (e.g. "://example.com/content.mov?ET=1209422976&CIP=10.0.0.1&KID=http://example.com/public/keys/123&DS=r:CFB03EDB33810AB6C79EE3C47FBD86D227D702F25F66C01CF03F59F1E005668D:s:57ED0E8DF7E786C87E39177DD3398A7FB010E6A4C0DC8AA71331A929A29EA24E")

3.2. Packaging the URI Signature

Apply the URI Signing Package Attribute by following the procedure below to generate the Signed URI.

1. Remove the Original URI portion from the message to obtain all the URI Signing Information Elements, including the URI signature (e.g. "ET=1209422976&CIP=10.0.0.1&KID=example:keys:123&&MD=da58513e8b309c1e8a9695baceba629d180b50b8").
2. Compute the URI Signing Package Attribute using Base-64 Data Encoding [RFC4648] on the message (e.g. "RVQ9MTIwOTQyMjk3NiZDSVA9MTAuMCAwLjEmS0lEPWV4YWlwbGU6a2V5czoxMjMmJk1EPWRhNTg1MTNlOGIzMd1jMWU4YTk2OTViYWNlYmE2MjlkMTgwYjUwYjg="). Note: This is the value for the URI Signing Package Attribute.

3. Copy the entire Original URI into a buffer to hold the message.
4. Check if the Original URI already contains a query string. If not, append a "?" character. If yes, append an "&" character.
5. Append the parameter name used to indicate the URI Signing Package Attribute, as communicated via the CDNI Metadata interface, followed by an "=". If none is communicated by the CDNI Metadata interface, it defaults to "URISigningPackage". For example, if the CDNI Metadata interface specifies "SIG", append the string "SIG=" to the message.
6. Append the URI Signing token to the message (e.g. "http://example.com/content.mov?URISigningPackage=RVQ9MTIwOTQyMjk3NiZDSVA9MTAuMC4wLjEmS0lEPWV4YW1wbGU6a2V5czoxMjMmJk1EPWRhNTg1MTNlOGIzMd1jMWU4YTk2OTViYWNlYmE2MjlkMTgwYjUwYjg="). Note: this is the completed Signed URI.

4. Validating a URI Signature

The process of validating a Signed URI can be divided into two sets of steps: validation of the information elements embedded in the Signed URI and validation of the URI Signature. Note it is possible to use some other algorithm and implementation as long as the same result is achieved.

4.1. Information element validation

Extract and validate the information elements embedded in the URI. Note that some steps are to be skipped if the corresponding URI Signing Information Element is not embedded in the Signed URI. The absence of a given Enforcement Information Element indicates enforcement of its purpose is not necessary in the CSP's distribution policy.

1. Extract the value from 'URISigningPackage' attribute. This value is the encoded URI Signing Package Attribute. If there are multiple instances of this attribute, the first one is used and the remaining ones are ignored. This ensures that the Signed URI can be validated despite a client appending another instance of the 'URISigningPackage' attribute.
2. Decode the string using Base-64 Data Encoding [RFC4648] (or another encoding method specified by configuration or CDNI metadata) to obtain all the URI Signing Information Elements (e.g. "ET=1209422976&CIP=10.0.0.1&KID=example:keys:123&&").

MD=da58513e8b309c1e8a9695baceba629d180b50b8").

3. Extract the value from "VER" if the information element exists in the query string. Determine the version of the URI Signing algorithm used to process the Signed URI. If the CDNI Metadata interface is used, check to see if the used version of the URI Signing algorithm is among the allowed set of URI Signing versions specified by the metadata. If this is not the case, the request is denied. If the attribute is not in the URI, then obtain the version number in another manner (e.g. configuration, CDNI metadata or default value).
4. Extract the value from "CIP" if the information element exists in the query string. Validate that the request came from the same IP address as indicated in the "CIP" attribute. If the IP address is incorrect, then the request is denied.
5. Extract the value from "ET" if the information element exists in the query string. Validate that the request arrived before expiration time based on the "ET" attribute. If the time expired, then the request is denied.
6. Extract the value from "MD" if the information element exists in the query string. The existence of this information element indicates a symmetric key is used.
7. Extract the value from "DS" if the information element exists in the query string. The existence of this information element indicates a asymmetric key is used.
8. If neither "MD" or "DS" attribute is in the URI, then no URI Signature exists and the request is denied. If both the "MD" and the "DS" information elements are present, the Signed URI is considered to be malformed and the request is denied.

4.2. Signature validation

Validate the URI Signature for the Signed URI.

1. Copy the Original URI, excluding the "scheme name" part, into a buffer to hold the message for performing the operations below.
2. Remove the "URISigningPackage" attribute from the message. Remove any subsequent part of the query string after the "URISigningPackage" attribute.
3. Append the decoded value from "URISigningPackage" attribute (which contains all the URI Signing Information Elements).

4. Depending on the type of key used to sign the URI, validate the message digest or digital signature for symmetric key or asymmetric keys, respectively.
 - A. For symmetric key, HMAC algorithm is used.
 - a. Extract the value from the "KID" information element, if it exists. Use the key identifier (e.g. "example:keys:123") to locate the shared key, which may be one of the keys available to use (i.e. set by configuration or CDNI metadata). If the information element is not in the URI Signing Package Attribute, then obtain the key in another manner (e.g. configuration or CDNI metadata). If the "KID" information element is present but its value is not in the allowable KID set as listed in the CDNI metadata, the request is denied.
 - b. Extract the value from the "HF" information element, if it exists. Determine the type of hash function (e.g. "MD5", "SHA-1", "SHA-256", "SHA-3") to use for HMAC. If the information element is not in the URI, the default hash function is SHA-1. If the "HF" information element is present but its value is not in the allowable "HF" set as listed in the CDNI metadata, the request is denied.
 - c. Extract the value from the "MD" information element. This is the received message digest.
 - d. Convert the message digest to binary format. This will be used to compare with the computed value later.
 - e. Remove the value part of the "MD" information element (but not the '=' character) from the message. The message is ready for validation of the message digest (e.g. "://example.com/content.mov?ET=1209422976&CIP=10.0.0.1&KID=example:keys:123&MD=").
 - f. Compute the message digest using the HMAC algorithm with the shared key and message as the two inputs to the hash function which is specified by the "HF" attribute.
 - g. Compare the result with the received message digest to validate the Signed URI.

- B. For asymmetric keys, a digital signature function is used.
- a. Extract the value from the "KID" information element, if it exists. Use the key identifier (e.g. "http://example.com/public/keys/123") to obtain the EC public key, which may be one of the keys available to use (i.e. set by configuration or CDNI metadata). If the information element is not in the URI, then obtain the key in another manner (e.g. configuration or CDNI metadata).
 - b. Extract the value from the "DSA" information element, if it exists. Determine the type of digital signature function (e.g. "RSA", "DSA", "EC-DSA") to use for calculating the Digital Signature. If the information element is not in the URI, the default digital signature function is EC-DSA. If the "DSA" information element is present but its value is not in the allowable "EC-DSA" set as listed in the CDNI metadata, the request is denied.
 - c. Extract the value from the "DS" information element. This is the digital signature.
 - d. Convert the digital signature to binary format. This will be used for verification later.
 - e. Remove the value part of the "DS" information element (but not the '=' character) from the message. The message is ready for validation of the digital signature (e.g. "://example.com/content.mov?ET=1209422976&CIP=10.0.0.1&KID=http://example.com/public/keys/123&DS=").
 - f. Compute the message digest using SHA-1 (without a key) for the message.
 - g. Verify the digital signature using the digital signature function (e.g. EC-DSA) with the public key, received digital signature, and message digest (obtained in previous step) as inputs. This validates the Signed URI.

5. Relationship with CDNI Interfaces

Some of the CDNI Interfaces need enhancements to support URI Signing. As an example: A Downstream CDN that supports URI Signing needs to be able to advertise this capability to the Upstream CDN. The Upstream

CDN needs to select a Downstream CDN based on such capability when the CSP requires access control to enforce its distribution policy via URI Signing. Also, the Upstream CDN needs to be able to distribute via the CDNI Metadata interface the information necessary to allow the Downstream CDN to validate a Signed URI . Events that pertain to URI Signing (e.g. request denial or delivery after access authorization) need to be included in the logs communicated through the CDNI Logging interface (Editor's Note: Is this within the scope of the CDNI Logging interface?).

5.1. CDNI Control Interface

URI Signing has no impact on this interface.

5.2. CDNI Footprint & Capabilities Advertisement Interface

The Downstream CDN advertises its capability to support URI Signing via the CDNI Footprint & Capabilities Advertisement interface (FCI). The supported version of URI Signing needs to be included to allow for future extensibility.

[Editor's Note: To be discussed with FCI authors]

5.3. CDNI Request Routing Redirection Interface

[Editor's Note: Debate the approach of dCDN providing the Signed URI vs. uCDN performing the signing function. List the pros/cons of each approach for the CDNI Request Routing Redirection interface (RI). Offer recommendation?]

The two approaches:

1. Downstream CDN provides the Signed URI
 - * Key distribution is not necessary
 - * Downstream CDN can use any scheme for Signed URI as long as the security level meets the CSP's expectation
2. Upstream CDN signs the URI
 - * Consistency with interative request routing method
 - * URI Signing works even when Downstream CDN does not have the signing function (which may be the case when the Downstream CDN operates only as a delivering CDN)

- * Upstream CDN can act as a conversion gateway for the requesting routing interface between Upstream CDN and CSP and request routing interface between Upstream CDN and Downstream CDN since these two interfaces may not be the same

5.4. CDNI Metadata Interface

The following CDNI Metadata objects are specified for URI Signing.

- o URI Signing enforcement flag. Specifically, this flag indicates if the access to content is subject to URI Signing. URI Signing requires the Downstream CDN to ensure that the URI must be signed and validated before content delivery. Otherwise, Downstream CDN does not perform validation regardless if URI is signed or not.
- o Designated key identifier used for URI Signing computation when the Signed URI does not contain the Key ID information element
- o Allowable Key ID set that the Signed URI's Key ID information element can reference
- o Designated hash function used for URI Signing computation when the Signed URI does not contain the Hash Function information element
- o Allowable Hash Function set that the Signed URI's Hash Function information element can reference
- o Designated digital signature function used for URI Signing computation when the Signed URI does not contain the Digital Signature Algorithm information element.
- o Allowable digital signature function set that the Signed URI's Digital Signature Algorithm information element can reference.
- o Designated version used for URI Signing computation when the Signed URI does not contain the VER attribute
- o Allowable version/algorithm set that the Signed URI's VER attribute can reference
- o Allowable set of Downstream CDNs that participate in URI Signing based on the symmetric key
- o Overwrite the default encoding method for URI Signing Attribute Set attribute? [Editor's Note: Do we need this?]
- o Overwrite the default name for the URL Signing Attribute Set attribute? [Editor's Note: Do we need this?]

Note that the Key ID information is not needed if only one key is provided by the CSP or the Upstream CDN for the content item or set of content items covered by the CDNI Metadata object. In the case of asymmetric keys, it's easy for any entity to sign the URI for content with a private key and provide the public key in the Signed URI. This just confirms that the URI Signer authorized the delivery. But it's necessary for the URI Signer to be the content owner. So, the CDNI Metadata interface MUST provide the public key for the content or information to authorize the received Key ID attribute.

5.5. CDNI Logging Interface

The Downstream CDN reports that enforcement of the access control was applied to the request for content delivery.

The following CDNI Logging field for URI Signing SHOULD be supported in the HTTP Request Logging Record as specified in CDNI Logging Interface [I-D.ietf-cdni-logging].

- o s-uri-signing:

- * format: 1DIGIT
- * field value: this characterises the uri signing validation performed by the Surrogate on the request. The allowed values are:
 - + "0" : no uri signature validation performed
 - + "1" : uri signature validation performed and validated
 - + "2" : uri signature validation performed and rejected
- * occurrence: there MUST be zero or exactly one instance of this field.

[Editor's note: Need to log these URI signature validation events (e.g. invalid client IP address, expired signed URI, incorrect URI signature, successful validation)?]

TBD: CDNI Logging interface is work in progress.

6. URI Signing Message Flow

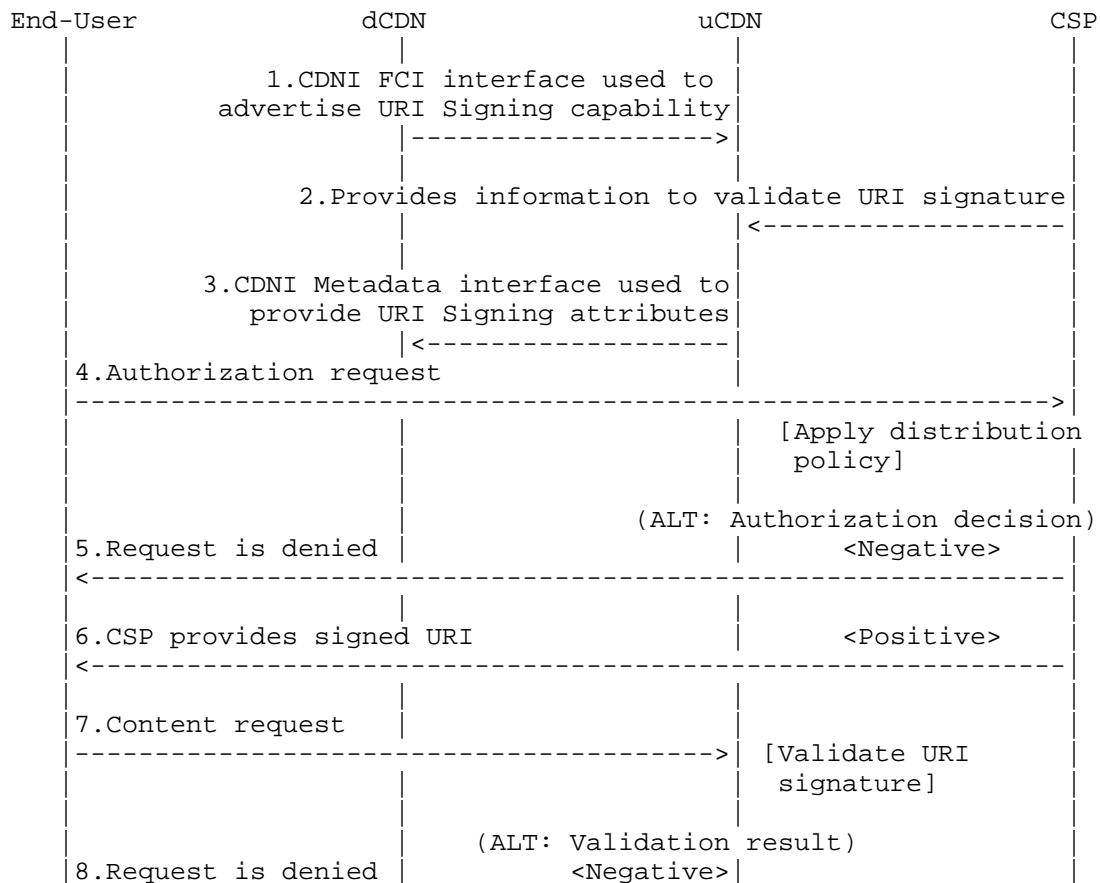
URI Signing supports both HTTP-based and DNS-based request routing. HMAC [RFC2104] defines a hash-based message authentication code allowing two parties that share a symmetric key or asymmetric keys to

establish the integrity and authenticity of a set of information (e.g. a message) through a cryptographic hash function.

6.1. HTTP Redirection

For HTTP-based request routing, HMAC is applied to a set of information that is unique to a given end user content request using key information that is specific to a pair of adjacent CDNI hops (e.g. between the CSP and the Authoritative CDN, between the Authoritative CDN and a Downstream CDN). This allows a CDNI hop to ascertain the authenticity of a given request received from a previous CDNI hop.

The URI signing scheme described below is based on the following steps (assuming HTTP redirection, iterative request routing and a CDN path with two CDNs). Note that Authoritative CDN and Upstream CDN are used exchangeably.



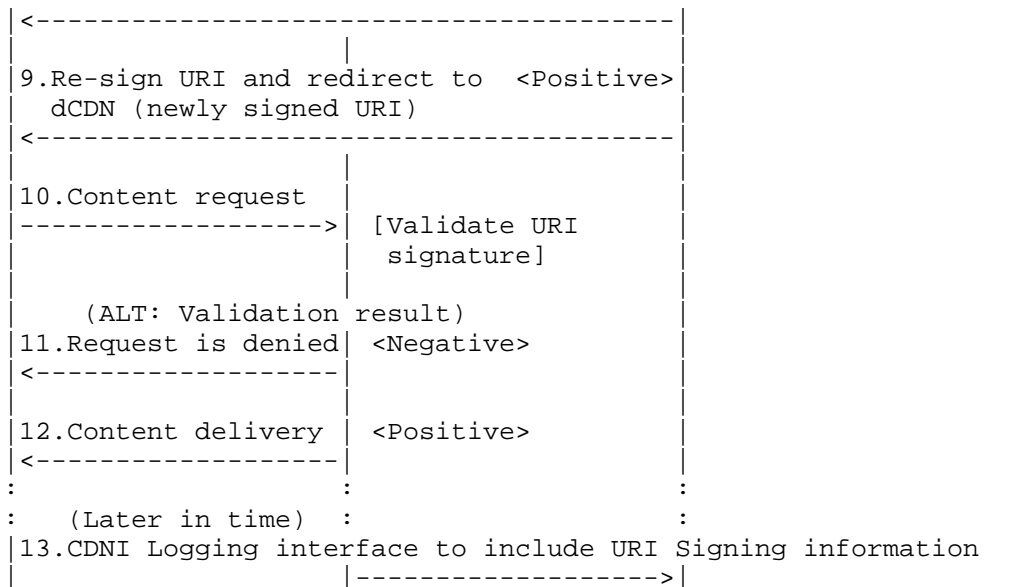


Figure 3: HTTP-based Request Routing with URI Signing

1. Using the CDNI Footprint & Capabilities Advertisement interface, the Downstream CDN advertises its capabilities including URI Signing support to the Authoritative CDN.
2. CSP provides to the Authoritative CDN the information needed to validate URI signatures from that CSP. For example, this information may include a hashing function, algorithm, and a key value.
3. Using the CDNI Metadata interface, the Authoritative CDN communicates to a Downstream CDN the information needed to validate URI signatures from the Authoritative CDN for the given CSP. For example, this information may include the URI query string parameter name for the URI Signing Package Attribute, a hashing algorithm and/or a key corresponding to the trust relationship between the Authoritative CDN and the Downstream CDN.
4. When a UA requests a piece of protected content from the CSP, the CSP makes a specific authorization decision for this unique request based on its arbitrary distribution policy
5. If the authorization decision is negative, the CSP rejects the request.

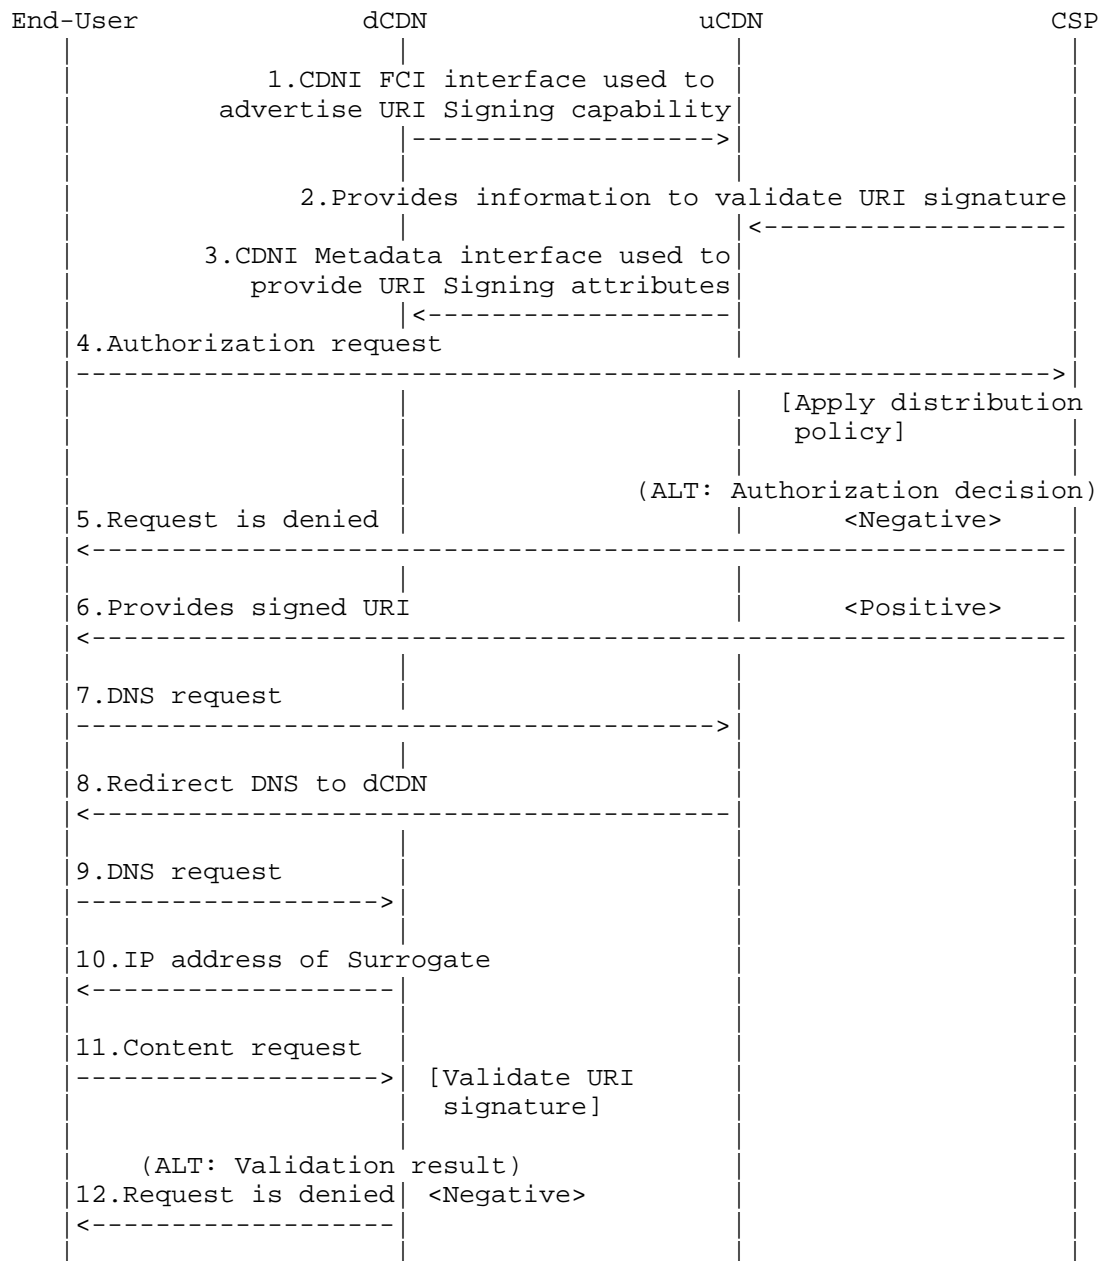
6. If the authorization decision is positive, the CSP computes a Signed URI that is based on unique parameters of that request and conveys it to the end user as the URI to use to request the content.
7. On receipt of the corresponding content request, the authoritative CDN validates the URI Signature in the URI using the information provided by the CSP.
8. If the validation is negative, the authoritative CDN rejects the request
9. If the validation is positive, the authoritative CDN computes a Signed URI that is based on unique parameters of that request and provides to the end user as the URI to use to further request the content from the Downstream CDN
10. On receipt of the corresponding content request, the Downstream CDN validates the URI Signature in the Signed URI using the information provided by the Authoritative CDN in the CDNI Metadata
11. If the validation is negative, the Downstream CDN rejects the request and sends an error code (e.g. 403) in the HTTP response.
12. If the validation is positive, the Downstream CDN serves the request and delivers the content.
13. At a later time, Downstream CDN reports logging events that includes URI signing information.

With HTTP-based request routing, URI Signing matches well the general chain of trust model of CDNI both with symmetric key and asymmetric keys because the key information only need to be specific to a pair of adjacent CDNI hops.

6.2. DNS Redirection

For DNS-based request routing, the CSP and Authoritative CDN must agree on a trust model appropriate to the security requirements of the CSP's particular content. Use of asymmetric public/private keys allows for unlimited distribution of the public key to Downstream CDNs. However, if a shared secret key is preferred, then the CSP may want to restrict the distribution of the key to a (possibly empty) subset of trusted Downstream CDNs. Authorized Delivery CDNs need to obtain the key information to validate the Signed UR, which is computed by the CSP based on its distribution policy.

The URI signing scheme described below is based on the following steps (assuming iterative DNS request routing and a CDN path with two CDNs). Note that Authoritative CDN and Upstream CDN are used exchangeably.



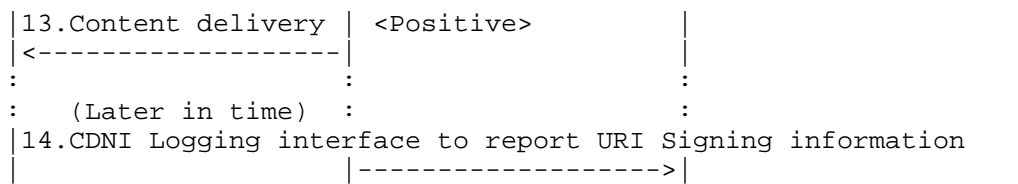


Figure 4: DNS-based Request Routing with URI Signing

1. Using the CDNI Footprint & Capabilities Advertisement interface, the Downstream CDN advertises its capabilities including URI Signing support to the Authoritative CDN.
2. CSP provides to the Authoritative CDN the information needed to validate cryptographic signatures from that CSP. For example, this information may include a hash function, algorithm, and a key.
3. Using the CDNI Metadata interface, the Authoritative CDN communicates to a Downstream CDN the information needed to validate cryptographic signatures from the CSP (e.g. the URI query string parameter name for the URI Signing Package Attribute). In the case of symmetric key, the Authoritative CDN checks if the Downstream CDN is allowed by CSP to obtain the shared secret key.
4. When a UA requests a piece of protected content from the CSP, the CSP makes a specific authorization decision for this unique request based on its arbitrary distribution policy.
5. If the authorization decision is negative, the CSP rejects the request
6. If the authorization decision is positive, the CSP computes a cryptographic signature that is based on unique parameters of that request and includes it in the URI provided to the end user to request the content.
7. End user sends DNS request to the authoritative CDN.
8. On receipt of the DNS request, the authoritative CDN redirects the request to the Downstream CDN.
9. End user sends DNS request to the Downstream CDN.
10. On receipt of the DNS request, the Downstream CDN responds with IP address of one of its Surrogates.

11. On receipt of the corresponding content request, the Downstream CDN validates the cryptographic signature in the URI using the information provided by the Authoritative CDN in the CDNI Metadata
12. If the validation is negative, the Downstream CDN rejects the request and sends an error code (e.g. 403) in the HTTP response.
13. If the validation is positive, the Downstream CDN serves the request and delivers the content.
14. At a later time, Downstream CDN reports logging events that includes URI signing information.

With DNS-based request routing, URI Signing matches well the general chain of trust model of CDNI when used with asymmetric keys because the only key information that need to be distributed across multiple CDNI hops including non-adjacent hops is the public key, that is generally not confidential.

With DNS-based request routing, URI Signing does not match well the general chain of trust model of CDNI when used with symmetric keys because the symmetric key information needs to be distributed across multiple CDNI hops including non-adjacent hops. This raises a security concern for applicability of URI Signing with symmetric keys in case of DNS-based inter-CDN request routing.

7. HTTP Adaptive Streaming

The authors note that in order to perform URI signing for individual content segments of HTTP Adaptive Bitrate content, specific URI signing mechanisms are needed. Such mechanisms are currently out-of-scope of this document. More details on this topic is covered in Models for HTTP-Adaptive-Streaming-Aware CDNI [RFC6983].

8. IANA Considerations

[Editor's note: (Is there a need to) register default value for URI Signing Package Attribute URI query string parameter name (i.e. URISigningPackage) to be used for URI Signing? Need anything from IANA?]

[Editor's note: To do: Convert to proper IANA Registry format]

This document requests IANA to create three new registries for the Information Elements and their defined values to be used for URI

Signing.

The following Enforcement Information Element names are allocated:

- o ET (Expiry time)
- o CIP (Client IP address)

The following Signature Computation Information Element names are allocated:

- o VER (Version): 1(Base)
- o KID (Key ID)
- o HF (Hash Function): "MD5", "SHA-1", "SHA-256", "SHA-3"
- o DSA (Digital Signature Algorithm): "RSA", "DSA", "EC-DSA"

The following URI Signature Information Element names are allocated:

- o MD (Message Digest)
- o DS (Digital Signature)

The IANA is requested to allocate a new entry to the CDNI Logging Field Names Registry as specified in CDNI Logging Interface [I-D.ietf-cdni-logging] in accordance to the "Specification Required" policy [RFC5226]

- o s-url-signing
- o [Editor's note: logging error codes are needed for URI Signing?]

The IANA is requested to allocate a new entry to the CDNI Metadata Field Names Registry as specified in CDNI Metadata Interface [I-D.ietf-cdni-metadata] in accordance to the "Specification Required" policy [RFC5226]

- o URI Signing Package URI query parameter name 1 Token
- o More metadata...

9. Security Considerations

This document describes the concept of URI Signing and how it can be used to provide access authorization in the case of interconnected

CDNs (CDNI). The primary goal of URI Signing is to make sure that only authorized UAs are able to access the content, with a Content Service Provider (CSP) being able to authorize every individual request. It should be noted that URI Signing is not a content protection scheme; if a CSP wants to protect the content itself, other mechanisms, such as DRM, are more appropriate.

In general, it holds that the level of protection against illegitimate access can be increased by including more Enforcement Information Elements in the URI. The current version of this document includes elements for enforcing Client IP Address and Expiration Time, however this list can be extended with other, more complex, attributes that are able to provide some form of protection against some of the vulnerabilities highlighted below.

That said, there are a number of aspects that limit the level of security offered by URI signing and that anybody implementing URI signing should be aware of.

Replay attacks: Any (valid) Signed URI can be used to perform replay attacks. The vulnerability to replay attacks can be reduced by picking a relatively short window for the Expiration Time attribute, although this is limited by the fact that any HTTP-based request needs a window of at least a couple of seconds to prevent any sudden network issues from preventing legitimate UAs access to the content. One way to reduce exposure to replay attacks is to include in the URI a unique one-time access ID. Whenever the Downstream CDN receives a request with a given unique access ID, it adds that access ID to the list of 'used' IDs. In the case an illegitimate UA tries to use the same URI through a replay attack, the Downstream CDN can deny the request based on the already-used access ID.

Illegitimate client behind a NAT: In cases where there are multiple users behind the same NAT, all users will have the same IP address from the point of view of the Downstream CDN. This results in the Downstream CDN not being able to distinguish between the different users based on Client IP Address and illegitimate users being able to access the content. One way to reduce exposure to this kind of attack is to not only check for Client IP but also for other attributes that can be found in the HTTP headers.

TBD: ...

The shared key between CSP and Authoritative CDN may be distributed to Downstream CDNs - including cascaded CDNs. Since this key can be used to legitimately sign a URL for content access authorization,

it's important to know the implications of a compromised shared key.

[Editor's note: Threat model cover in the Security section - Prevent client from spoofing URI (Ray) - Security implications - The scope of protection by URI Signing - Protects against DoS (network bandwidth and other nodes besides the edge cache); limits the time window.]

10. Privacy

The privacy protection concerns described in CDNI Logging Interface [I-D.ietf-cdni-logging] apply when the client's IP address (CIP attribute) is embedded in the Signed URI. This means that, when anonymization is enabled, the URI Signing Package Attribute MUST be removed from the logging record.

11. Acknowledgements

The authors would like to thank the following people for their contributions in reviewing this document and providing feedback: Kevin Ma, Ben Niven-Jenkins, Thierry Magnien, Dan York, Bhaskar Bhupalam, Matt Caulfield, and Samuel Rajakumar .

12. References

12.1. Normative References

- [I-D.ietf-cdni-logging]
Faucheur, F., Bertrand, G., Oprescu, I., and R. Peterkofsky, "CDNI Logging Interface", draft-ietf-cdni-logging-10 (work in progress), March 2014.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [RFC6707] Niven-Jenkins, B., Le Faucheur, F., and N. Bitar, "Content Distribution Network Interconnection (CDNI) Problem Statement", RFC 6707, September 2012.

12.2. Informative References

- [I-D.ietf-cdni-framework]
Peterson, L., Davie, B., and R. Brandenburg, "Framework for CDN Interconnection", draft-ietf-cdni-framework-10 (work in progress), March 2014.

- [I-D.ietf-cdni-metadata]
Niven-Jenkins, B., Murray, R., Watson, G., Caulfield, M.,
Leung, K., and K. Ma, "CDN Interconnect Metadata",
draft-ietf-cdni-metadata-06 (work in progress),
February 2014.
- [I-D.ietf-cdni-requirements]
Leung, K. and Y. Lee, "Content Distribution Network
Interconnection (CDNI) Requirements",
draft-ietf-cdni-requirements-17 (work in progress),
January 2014.
- [RFC2104] Krawczyk, H., Bellare, M., and R. Canetti, "HMAC: Keyed-
Hashing for Message Authentication", RFC 2104,
February 1997.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform
Resource Identifier (URI): Generic Syntax", STD 66,
RFC 3986, January 2005.
- [RFC4648] Josefsson, S., "The Base16, Base32, and Base64 Data
Encodings", RFC 4648, October 2006.
- [RFC5952] Kawamura, S. and M. Kawashima, "A Recommendation for IPv6
Address Text Representation", RFC 5952, August 2010.
- [RFC6770] Bertrand, G., Stephan, E., Burbridge, T., Eardley, P., Ma,
K., and G. Watson, "Use Cases for Content Delivery Network
Interconnection", RFC 6770, November 2012.
- [RFC6983] van Brandenburg, R., van Deventer, O., Le Faucheur, F.,
and K. Leung, "Models for HTTP-Adaptive-Streaming-Aware
Content Distribution Network Interconnection (CDNI)",
RFC 6983, July 2013.

Authors' Addresses

Kent Leung
Cisco Systems
3625 Cisco Way
San Jose 95134
USA

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

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

Bill Downey
Verizon Labs
60 Sylvan Road
Waltham, Massachusetts 02451
USA

Phone: +1 781 466 2475
Email: william.s.downey@verizon.com

Ray van Brandenburg
TNO
Brassersplein 2
Delft, 2612CT
the Netherlands

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

Scott Leibrand
Limelight Networks
222 S Mill Ave
Tempe, AZ 85281
USA

Phone: +1 360 419 5185
Email: sleibrand@llnw.com

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 24, 2014

K. Ma
Azuki Systems, Inc.
October 21, 2013

CDNI Footprint & Capabilities Advertisement Interface
draft-ma-cdni-capabilities-04

Abstract

Content Distribution Network Interconnection (CDNI) is predicated on the ability of downstream CDNs (dCDNs) to handle end-user requests in a functionally equivalent manner to the upstream CDN (uCDN). The uCDN must be able to assess the ability of the dCDN to handle individual requests. The CDNI Footprint & Capabilities Advertisement interface (FCI) is provided for the advertisement of capabilities and the footprints to which they apply by the dCDN to the uCDN. This document describes an approach to implementing the CDNI FCI.

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 April 24, 2014.

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	2
1.1. Terminology	4
2. Capabilities	4
2.1. Delivery Protocol	6
2.2. Acquisition Protocol	7
2.3. Redirection Mode	9
2.4. Logging Capabilities	10
2.5. Metadata Capabilities	10
3. Capability Advertisement	10
3.1. Capability Initialization	11
3.2. Capability Reset	11
4. IANA Considerations	12
5. Security Considerations	12
6. Acknowledgements	12
7. References	12
7.1. Normative References	12
7.2. Informative References	13
Appendix A. Capability Aggregation	13
A.1. Downstream CDN Aggregation	13
A.2. Internal Request Router Aggregation	15
A.3. Internal Capability Aggregation	17
Author's Address	18

1. Introduction

The need for footprint and capabilities advertisement in CDNI is described in the CDNI requirements document [I-D.ietf-cdni-requirements]. Requirements FCI-1 and FCI-2 describe the need to allow dCDNs to communicate capabilities to the uCDN. Requirement FCI-3 describes how a uCDN may aggregate the footprint and capabilities information for all cascaded dCDNs and use the

aggregated information in advertisements to CDNs further upstream. This concept of aggregation can apply to both organizationally different dCDNs (e.g., other CDN providers, or different business units within a larger organization) or logical entities within the same CDN (e.g., using multiple request routers for scalability reasons, to segregate surrogates based on specific protocol support, or to segregate surrogates based on software version or feature level, etc.).

Appendix A contains more detailed descriptions of different footprint and capabilities management scenarios, but it is important to note that it is the ability of the dCDN to service each request in a functionally equivalent manner as the uCDN that is important, not the physical layout of resources through which it services the request. The aggregation of resource knowledge by the dCDN into a simple set of capabilities and their affective footprints, that is then advertised to the uCDN enables efficient decision making at each delegation point in the CDN interconnection hierarchy.

It is assumed that an authoritative request router in each CDN will be responsible for aggregating and advertising capabilities information in a dCDN, and receiving and aggregating capabilities information in the uCDN. The CDNI Footprint & Capabilities Advertisement interface (FCI) along with the CDNI Request Routing Redirection interface (RI) make up the CDNI Request Routing Interface. As there is no other centralized CDNI controller, the authoritative request router seems the most logical place for capabilities aggregation to occur, as it is the request router that needs such information to make delegation decisions. The protocol defined herein may be implemented as part of an entity other than an authoritative request router, but for the purposes of this discussion, the authoritative request router is assumed to be the centralized capabilities aggregation point.

Though there is an obvious need for the ability to exchange and update footprint and capability information in real-time, it is assumed that capabilities do not change very often. It is also assumed that the capabilities are not by themselves useful for making delegation decisions. Capability information is assumed to be input into business logic. It is the business logic which provides the algorithms for delegation decision making. The definition of business logic occurs outside the scope of CDNI and outside the timescale of footprint and capability advertisement. It may be the case that the business logic anticipates and reacts to changes in dCDN capabilities. However, it may also be the case that business logic is tailored through offline processes as dCDN capabilities change. The FCI is agnostic to the business processes employed by any given uCDN. The footprints and capabilities that are advertised

over the FCI may be used by the uCDN at its discretion to implement delegation rules. Setting proper defaults in the business logic should prevent any unwanted delegation from occurring when dCDN capabilities change, however, that is beyond the scope of this discussion.

1.1. Terminology

This document uses the terminology defined in section 1.1 of the CDNI Framework [I-D.ietf-cdni-framework] document.

2. Capabilities

As described in Requirement FCI-2, there is a basic set of capabilities that must be supported by the FCI for the uCDN to be able to determine if the dCDN is functionally able to handle a given request. The CDNI Footprint and Capabilities Semantics [I-D.ietf-cdni-footprint-capabilities-semantics] document lists mandatory capabilities types:

- o Delivery Protocol
- o Acquisition Protocol
- o Redirection Mode
- o CDNI Logging Capabilities
- o CDNI Metadata Capabilities

The following sections describe each of the capabilities in further detail, however, all of the capabilities can be described using the same general format. A capability object is specified as:

CapabilityObject:

 Name: Identifier for the capability

 Values: List of supported options for the capability

 Footprint: Optional list of footprint objects

The list of valid capability options for a given capability will be specific to the given capability type. Though the degenerate case may exist where the range of option values is a single value, it is anticipated that all capability types will have more than one capability option value. For consistency in the model, all capability types are implemented with lists of values. To optimize actions on the entire range of capability option values for a given capability type, the capability option value "ALL" is reserved and

MUST be supported by all capability types. For completeness, the capability option value "NONE" is also reserved and MUST be supported by all capability types. If a reserved value is specified, it MUST be the only entry in the capability value list.

The footprint restriction list is optional. When included, the footprint restriction list contains a list of generic footprint objects. The footprint object is specified as:

FootprintObject:

Type: Identifier for the capability

Values: List of footprint entries

Mode: Optional footprint application directive

The footprint type specifies the format of the footprint value entries. The footprint object type field contains a registered footprint type value from the "CDNI Metadata Footprint Types" registry. The CDNI Footprint and Capabilities Semantics [I-D.ietf-cdni-footprint-capabilities-semantics] document lists the mandatory footprint types as: ISO Country Code, AS number, and IP-prefix. The CDNI Metadata Interface [I-D.ietf-cdni-metadata] document defines the footprint type registry and the initial values for the mandatory footprint types. It also describes the process for registering additional optional footprint types. The footprint value "GLOBAL" is reserved and MUST be supported by all footprint types. If the reserved value "GLOBAL" is specified, it MUST be the only entry in the footprint value list.

The optional footprint mode describes how the footprint should be applied. The valid footprint mode values are: "replace", "include", and "exclude". The footprint mode "replace" (represented by the integer value 0) indicates that all previous footprint information for the given footprint type (and for the capabilities to which the current footprint object applies) MUST be replaced in its entirety with the footprint information specified in the accompanying footprint value list. The footprint mode "include" (represented by the integer value 1) indicates that the any existing footprint information for the given footprint type (and for the capabilities to which the current footprint object applies) SHOULD be augmented with the footprint information specified in the accompanying footprint value list. The footprint mode "exclude" (represented by the integer value 2) indicates that the any existing footprint information for the given footprint type (and for the capabilities to which the current footprint object applies) SHOULD be reduced by the footprint information specified in the accompanying footprint value list. If no footprint mode value is specified, the default mode SHALL be understood to be "replace".

The footprint restriction list MUST NOT contain multiple footprint objects of the same type. Footprint restriction information MAY be specified using multiple different footprint types. If no footprint restriction list is specified (or an empty list is specified), it SHALL be understood that all footprint types MUST be reset to "GLOBAL" coverage.

Note: Further optimization of the footprint object to provide quality information for a given footprint is certainly possible, however, it is not critical to the basic interconnection of CDNs. The ability to transfer quality information in capabilities advertisements may be desirable and is noted here for completeness, however, the specifics of such mechanisms are outside the scope of this document.

Multiple capability objects of the same capability type are allowed within a given FCI message as long as the capability option values do not overlap, i.e., a given capability option value MUST NOT show up in multiple capability objects within a single FCI message. If multiple capability objects for a given capability type exist, those capability objects SHOULD have different footprint restrictions; capability objects of a given capability type with identical footprint restrictions SHOULD be combined into a single capability object.

2.1. Delivery Protocol

The delivery protocol refers to the protocol over which an end user (EU) has requested content. If a dCDN does not support the protocol requested by the client, then the dCDN is not a viable candidate for delegation.

Though the delivery protocol is specified in the URI scheme (as defined in RFC3986 [RFC3986]) of the client request URL, protocol feature subsets or augmented protocol feature sets MAY be defined and SHOULD correspond with the protocols supported by the ProtocolACL defined in the CDNI Metadata Interface [I-D.ietf-cdni-metadata] document. The CDNI Metadata Interface document defines the "CDNI Metadata Protocols" registry and the initial supported protocol values. It also describes the process for registering additional protocols.

The delivery protocol capability object MUST support a list of protocols for a given footprint. The delivery protocol capability SHOULD support optional footprint restriction information. The following example shows two lists of protocols with different footprints.

```

{
  "capabilities": [
    { "name": "delivery_protocol",
      "values": [
        "HTTP",
        "RTSP",
        "MMS"
      ]
    },
    { "name": "delivery_protocol",
      "values": [
        "RTMP",
        "HTTPS"
      ],
      "footprint": [
        { "type": "IPv4",
          "values": [
            "10.1.0.0/16",
            "10.10.10.0/24"
          ]
        }
      ]
    }
  ]
}

```

In the above example, the three protocols HTTP, RTSP, and MMS are supported globally, while the protocols RTMP and HTTPS are only supported in a restricted footprint (in this case, specified by IP-prefix).

A given protocol **MUST NOT** appear in multiple capability object value lists, within a given FCI message.

[Ed. need to add reference to registry where the protocol values are defined, once they are finalized in the semantics/metadata draft.]

2.2. Acquisition Protocol

The acquisition protocol refers to the protocol over which the dCDN may acquire content from the uCDN. If a dCDN does not support any of the protocols offered by the uCDN, then the dCDN is not a viable candidate for delegation.

Though the acquisition protocol is disseminated to the dCDN in the URI scheme (as defined in RFC3986 [RFC3986]) of the URL provided by the uCDN via the CDNI Metadata Interface [I-D.ietf-cdni-metadata],

protocol feature subsets or augmented protocol feature sets MAY be defined and SHOULD correspond with the protocols supported by the ProtocolACL defined in the CDNI Metadata Interface [I-D.ietf-cdni-metadata] document. The CDNI Metadata Interface document defines the "CDNI Metadata Protocols" registry and the initial supported protocol values. It also describes the process for registering additional protocols.

The acquisition protocol capability object MUST support a list of protocols for a given footprint. The acquisition protocol capability SHOULD support optional footprint restriction information. The following example shows two lists of protocols with different footprints.

```
{
  "capabilities": [
    { "name": "acquisition_protocol",
      "values": [
        "HTTP",
        "FTP"
      ]
    },
    { "name": "acquisition_protocol",
      "values": [
        "SFTP",
        "HTTPS"
      ],
      "footprint": [
        { "type": "ASN",
          "values": [
            "0",
            "65535"
          ],
          "mode": 2
        }
      ]
    }
  ]
}
```

In the above example, the two protocols HTTP and FTP are supported globally, while the protocols SFTP and HTTPS are updated to only be supported in a reduced restricted footprint (in this case, specified by ASN).

A given protocol MUST NOT appear in multiple capability object value lists, within a given FCI message.

2.3. Redirection Mode

The redirection mode refers to the method(s) employed by request routers to perform request redirection. The CDNI framework [I-D.ietf-cdni-framework] document describes four possible request routing modes:

- o DNS iterative (DNS-I)
- o DNS recursive (DNS-R)
- o HTTP iterative (HTTP-I)
- o HTTP recursive (HTTP-R)

The CDNI Footprint and Capabilities Semantics [I-D.ietf-cdni-footprint-capabilities-semantics] defines the "CDNI Capabilities Redirection Modes" registry and the initial supported redirection mode values shown in parentheses above. It also describes the process for registering additional redirection modes.

If a dCDN supports only a specific mode or subset of modes that does not overlap with the modes supported by the uCDN, then the dCDN is not a viable candidate for delegation.

The redirection mode capability object MUST support a list of redirection modes for a given footprint. The redirection mode capability SHOULD support optional footprint restriction information. The following XML-encoded example shows two lists of modes with different footprints.

```
{
  "capabilities": [
    { "name": "redirection_mode",
      "values": [
        "DNS-I",
        "HTTP-I"
      ]
    },
    { "name": "redirection_mode",
      "values": [
        "DNS-R",
        "HTTP-R"
      ]
    },
    "footprint": [
      { "type": "ASN",
        "values": [
          "9"
        ]
      }
    ]
  }
}
```

```
    ],
    "mode": 0
  },
  {
    "type": "IPv6",
    "values": [
      "8765:4321::/36"
    ]
  }
]
}
```

In the above example, iterative redirection is supported globally, while recursive redirection is only supported in a restricted footprint (in this case, specified by both ASN and IP-prefix).

A given mode **MUST NOT** appear in multiple capability object value lists, within a given FCI message.

2.4. Logging Capabilities

The CDNI Logging interface [I-D.ietf-cdni-logging] document describes optional logging fields and functionality which may be optional for a dCDN to implement. If a dCDN does not support certain logging parameters which may affect billing agreements or legal requirements of the uCDN, then the dCDN is not a viable candidate for delegation.

[Ed. need to update this section once the list of logging capabilities is finalized in the semantics/logging draft.]

2.5. Metadata Capabilities

The CDNI Metadata interface [I-D.ietf-cdni-metadata] document describes generic metadata types which may be optional for a dCDN to implement, but which, if present, are mandatory-to-enforce. If a dCDN does not support certain metadata types which are designated mandatory-to-enforce and may affect the correctness or security of the content being delivered, then the dCDN is not a viable candidate for delegation.

[Ed. need to update this section once the list of metadata capabilities is finalized in the semantics/metadata draft.]

3. Capability Advertisement

The FCI relies an HTTP-based protocol using the GET and POST methods. The uCDN request router may at any time query the dCDN request router for the fully capability set of the dCDN. In addition, the dCDN request router SHOULD asynchronously issue HTTP POSTs of FCI messages to the uCDN request router (see Appendix A for detailed descriptions of authoritative request router capabilities aggregation scenarios) whenever changes in the its already advertised capabilities occur. It is assumed that the dCDN request router has been configured with the uCDN request router address, and the uCDN request router has been configured with the dCDN request router address, either through the CDNI Control interface (CI) bootstrapping function, or through some other out-of-band configuration.

3.1. Capability Initialization

In lieu of any out-of-band pre-configured capability information, when the FCI is first brought up between a uCDN and dCDN, the uCDN SHOULD assume that the dCDN has no CDNI capabilities. If an out-of-band capability baseline has been exchanged, the uCDN MAY use that information to initialize its capabilities database. In either case, the uCDN SHOULD verify the initial state of the dCDN (as a temporary outage may be affecting availability in the dCDN).

The dCDN MUST support sending its entire set of capabilities to the uCDN using the footprint mode "replace", in order to (re)initialize the uCDN capabilities database. In response to an FCI GET request from the uCDN, the dCDN MUST send its entire set of capabilities to the uCDN using the footprint mode "replace".

[Ed.: The alternative to using a pull from the uCDN is to use the triggers interface for a triggered push, however, this would not be triggering a CDN function, it would be triggering an FCI function, so given that there is no asynchronous action required by the dCDN, it seems that reducing inter-dependency on other CDNI interfaces makes the most sense in this case.]

3.2. Capability Reset

When using the footprint modes "include" and/or "exclude" for partial update of the footprint for a given capability, there is a dependency and expectation as to what the uCDN believes the current footprint of that capability to be. If any in the sequence of footprint update messages are lost, there could be a loss of coherency between the uCDN and the dCDN. To help prevent this situation, each FCI POST MUST include a sequence number if it intends to use the footprint modes "include" and/or "exclude". The sequence number SHALL be represented as a 32 bit unsigned integer that wraps. The sequence number MUST appear in the "CDNI-FCI-seq" HTTP header and be incremented for every FCI message sent by the dCDN to a given uCDN.

If the uCDN ever receives an out of sequence FCI message from a given dCDN, the uCDN SHOULD reinitialize its capabilities database using an FCI GET request.

4. IANA Considerations

This memo includes no request to IANA.

5. Security Considerations

There are a number of security concerns associated with the FCI. The FCI essentially provides configuration information which the uCDN uses to make request routing decisions. Injection of fake capability advertisement messages or the interception and discard of real capability advertisement messages may be used for denial of service (e.g., by falsely advertising or deleting capabilities or preventing capability advertisements from reaching the uCDN). dCDN capability advertisements MUST be authenticated by the uCDN to prevent unauthorized FCI message injection. uCDN FCI servers MUST be authenticated by the dCDN to prevent unauthorized interception of FCI messages. TLS with client authentication SHOULD be used for all FCI implementations. Deployments in controlled environments where physical security and IP address white-listing is employed MAY choose not to use TLS.

6. Acknowledgements

The authors would like to thank Ray van Brandenburg, Gilles Bertrand, and Scott Wainner for their timely reviews and invaluable comments.

7. References

7.1. Normative References

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

- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005.

7.2. Informative References

- [I-D.ietf-cdni-footprint-capabilities-semantics]
Seedorf, J., Peterson, J., Previdi, S., Brandenburg, R.,
and K. Ma, "CDNI Request Routing: Footprint and
Capabilities Semantics", draft-ietf-cdni-footprint-
capabilities-semantics-00 (work in progress), July 2013.
- [I-D.ietf-cdni-framework]
Peterson, L. and B. Davie, "Framework for CDN
Interconnection", draft-ietf-cdni-framework-06 (work in
progress), October 2013.
- [I-D.ietf-cdni-logging]
Faucheur, F., Bertrand, G., Oprescu, I., and R.
Peterkofsky, "CDNI Logging Interface", draft-ietf-cdni-
logging-08 (work in progress), October 2013.
- [I-D.ietf-cdni-metadata]
Niven-Jenkins, B., Murray, R., Watson, G., Caulfield, M.,
Leung, K., and K. Ma, "CDN Interconnect Metadata", draft-
ietf-cdni-metadata-02 (work in progress), July 2013.
- [I-D.ietf-cdni-requirements]
Leung, K. and Y. Lee, "Content Distribution Network
Interconnection (CDNI) Requirements", draft-ietf-cdni-
requirements-11 (work in progress), October 2013.

Appendix A. Capability Aggregation

The following sections show examples of three aggregation scenarios. In each case, CDN-U is the ultimate uCDN and CDN-P is the penultimate CDN which must perform capabilities aggregation.

A.1. Downstream CDN Aggregation

Figure A1 shows five organizationally different CDNs: CDN-U, CDN-P, and CDNs A, B, and C, the dCDNs of CDN-P which are being aggregated. Given the setup shown in Figure A1, we can construct a number of use cases, based on the coverage areas of each dCDN (i.e., CDNs P, A, B, and C). Note: In all cases, the reachability of the uCDN (i.e., CDN-U) is a don't care as it is assumed that the uCDN knows its own coverage area and is likely to favor itself in most situations, and if it has decided that it needs to delegate to a dCDN, then the only relevant question is if the dCDN can handle the request.

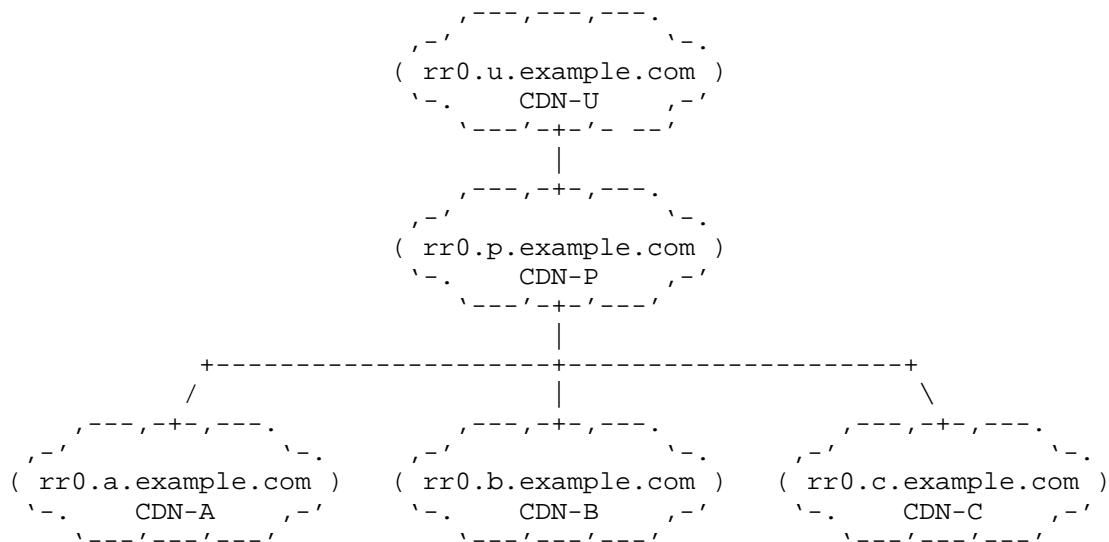


Figure A1: CDNI dCDN Request Router Aggregation

- o None of the four dCDNs (CDNs P, A, B, and C) have global reachability. In this case, each CDN is likely to advertise footprint information with its capabilities, specifying its reachability. When CDN-P advertises capabilities to CDN-U, it may advertise the aggregate footprint of itself and CDNs A, B, and C. Note: CDN-P MAY exclude any dCDN, and consequently its footprint, per its own internal aggregation decision criteria.
- o All four dCDNs (CDNs P, A, B, and C) have global reachability. In this case, none of the CDNs is likely to advertise any footprint information as none have any footprint restrictions. When CDN-P advertises capabilities to CDN-U, the aggregate of all global reachability is global reachability.
- o Some of the four dCDNs (CDNs P, A, B, and C) have global reachability and some do not. In this case, even though some

dCDNs do not have global reachability, the aggregate of some dCDNs having global reachability and some not should still be global reachability (for the given capability). When CDN-P advertises capabilities to CDN-U, CDN-P may advertise capabilities for which at least one dCDN has global reach as being supported with global reachability. It is up to the CDN-P request router to properly select a dCDN to process individual client requests and not choose a dCDN whose restricted footprint makes it unsuitable for delivering the requested content.

A.2. Internal Request Router Aggregation

Figure A2 shows CDN-U and CDN-P where CDN-P internally has four request routers: the authoritative request router rr0, and three other request routers rr1, rr2, and rr3. The use of multiple request routers may be used to distribute request routing load across resources, possibly in different geographic regions covered by CDN-P. Similar to Figure A1, the setup shown in Figure A2 requires the authoritative request router rr0 in CDN-P to aggregate capabilities information from downstream request routers rr1, rr2, and rr3. The primary difference between the scenario is that the request routers in Figure A2 are logically within the same CDN-P organization. The same reachability scenarios apply to Figure A2 as with Figure A1.




Figure A2: Local CDN Request Router Aggregation


- o None of the four CDN-P request routers have global reachability. In this case, each request router is likely to advertise footprint information with its capabilities, specifying its reachability. When rr0 advertises capabilities to CDN-U, it may advertise the aggregate footprint of itself and rr1, rr2, and rr3.
- o All four CDN-P request routers have global reachability. In this case, none of the request routers is likely to advertise any footprint information as none has any footprint restrictions. When rr0 advertises capabilities to CDN-U, the aggregate of all global reachability is global reachability.
- o Some of the four CDN-P request routers have global reachability and some do not. In this case, even though some request routers do not have global reachability, the aggregate of some request routers having global reachability and some not should still be global reachability (for the given capability). When rr0 advertises capabilities to CDN-U, CDN-P may advertise capabilities for which at least one request router has global reach as being supported with global reachability. It is up to the authoritative request router rr0 to properly select from the other request routers for any given request, and not choose a request router

whose restricted footprint makes it unsuitable for delivering the requested content.

A.3. Internal Capability Aggregation

Figure A3 shows CDN-U and CDN-P where the delivery network of CDN-P is segregated by delivery protocol (e.g., RTSP, HTTP, and RTMP). Figure A3 differs from Figures A1 and A2 in that request router rr0 of CDN-P is not aggregating the capabilities advertisements of multiple other downstream request routers, but rather it is managing the disparate capabilities across resources within its own local CDN. Though not every delivery node has the same protocol capabilities, the aggregate delivery protocol capabilities advertised by CDN-A may include all delivery protocols. Note, Figure A3 should not be construed to imply anything about the coverage areas for each delivery protocol. They may all support the same delivery footprint, or they may have different delivery footprints. It is the responsibility of the request router rr0 to properly assign protocol-appropriate delivery nodes to individual content requests. If certain protocols have limited reachability, CDN-P may advertise footprint restrictions for each protocol.

It should be noted that though the delivery protocol capability was selected for this example, the concept of internal capability aggregation applies to all capabilities as discussed below.



\\ _ _ _ _ / _ _ _ _ / _ _ _ _ / _ _ _ _ / _ _ _ _ / _ _ _ _ / _ _ _ _ / _ _ _ _ /

Figure A3: Local CDN Capability Segregation

Another situation in which physical footprint may not matter in an aggregated view has to do with feature support (e.g., new CDNI metadata features or new redirection modes). Situations often arise when phased roll-out of software upgrades, or staging network segregation result in only certain portions of a CDN's resources supporting the new feature set. The dCDN has a few options in this case:

- o Enforce atomic update: The dCDN does not advertise support for the new capability until all resources have been upgraded to support the new capability.
- o Transparent segregation: The dCDN advertises support for the new capability, and when requests are received that require the new capability, the dCDN request router properly selects a resource which supports that capability.
- o Advertised segregation: The dCDN advertises support for the new capability with a footprint restriction allowing the uCDN to make delegation decisions based on the dCDN's limit support.

The level of aggregation employed by the dCDN is likely to vary as business relationships dictate, however, the FCI should support all possible modes of operation.

Author's Address

Kevin J. Ma
Azuki Systems, Inc.
43 Nagog Park
Acton, MA 01720
USA

Phone: +1 978-844-5100
Email: kevin.ma@azukisystems.com

CDNI
Internet-Draft
Intended status: Informational
Expires: April 24, 2014

J. Seedorf
NEC
Y. Yang
Yale
October 21, 2013

CDNI Footprint and Capabilities Advertisement using ALTO
draft-seedorf-cdni-request-routing-alto-05

Abstract

Network Service Providers (NSPs) are currently considering to deploy Content Delivery Networks (CDNs) within their networks. As a consequence of this development, there is a need for interconnecting these local CDNs. The necessary interfaces for inter-connecting CDNs are currently being defined in the Content Delivery Networks Interconnection (CDNI) WG. This document focuses on the CDNI Footprint & Capabilities Advertisement interface (FCI). Specifically, this document outlines how the solutions currently being defined in the Application Layer Traffic Optimization (ALTO) WG can facilitate Footprint & Capabilities Advertisement in a CDNI context, i.e. how the CDNI FCI can be realised with the ALTO protocol. Concrete examples of how ALTO can be integrated within CDNI request routing and in particular in the process of selecting a downstream CDN are given. The examples in this document are based on the use cases and examples currently being discussed in the CDNI WG.

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 April 24, 2014.

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	2
2. ALTO within CDNI Request Routing	3
3. Assumptions and High-Level Design Considerations	4
3.1. General Assumptions and Consideration	4
3.2. Semantics for Footprint/Capabilities Advertisement	5
4. Selection of a Downstream CDN with ALTO	7
4.1. Footprint and Capabilities Advertisement using ALTO Network Map and PID Properties	7
4.2. Conveying additional information with ALTO Cost Maps	8
4.3. Example of Selecting a Downstream CDN based on ALTO Maps	9
4.4. Advantages of using ALTO	10
5. Useful ALTO extensions for CDNI Request Routing	10
6. Security Considerations	12
7. Summary and Outlook	12
8. Acknowledgements	12
9. References	13
9.1. Normative References	13
9.2. Informative References	13
Authors' Addresses	15

1. Introduction

Many Network Service Providers (NSPs) are currently considering or have already started to deploy Content Delivery Networks (CDNs) within their networks. As a consequence of this development, there is a need for interconnecting these local CDNs. Content Delivery Networks Interconnection (CDNI) has the goal of standardizing protocols to enable such interconnection of CDNs [RFC6707].

The CDNI problem statement [RFC6707] envisions four interfaces to be standardized within the IETF for CDN interconnection:

- o CDNI Request Routing Interface

- o CDNI Metadata Interface
- o CDNI Logging Interface
- o CDNI Control Interface

This document focuses solely on the CDNI Request Routing Interface, which can be further divided into two interfaces (see [RFC6707] for a detailed description): the CDNI Request Routing Redirection interface (RI), and the CDNI Footprint & Capabilities Advertisement interface (FCI). This document presents how one may use ALTO as a protocol for CDNI Footprint & Capabilities Advertisement. Concrete examples of how the CDNI FCI can be implemented with the ALTO protocol [I-D.ietf-alto-protocol] are given. The examples used in this document are based on the use cases and request routing proposals currently being discussed in the CDNI WG [RFC6770] [I-D.peterson-CDNI-strawman] and in the ALTO WG [I-D.jenkins-alto-cdn-use-cases].

A previous version of this document [I-D.seedorf-alto-for-cdni] contained detailed examples of actual request routing and surrogate selection with ALTO, i.e. how ALTO could be used for implementing the CDNI Request Routing Redirection interface (RI). This version solely focuses on implementing the CDNI Footprint & Capabilities Advertisement interface (FCI) with ALTO, i.e. the selection of a downstream CDN and how ALTO can support such downstream CDN selection.

Throughout this document, we use the terminology for CDNI defined in [I-D.ietf-cdni-problem-statement].

2. ALTO within CDNI Request Routing

The main purpose of the CDNI Request Routing Interface is described in [RFC6707] as follows: "The CDNI Request Routing interface enables a Request Routing function in an Upstream CDN to query a Request Routing function in a Downstream CDN to determine if the Downstream CDN is able (and willing) to accept the delegated Content Request. It also allows the Downstream CDN to control what should be returned to the User Agent in the redirection message by the upstream Request Routing function." On a high level, the scope of the CDNI Request Routing Interface therefore contains two main tasks:

- o A) Determining if the downstream CDN is willing to accept a delegated content request
- o B) Redirecting the content request coming from an upstream CDN to the proper entry point or entity in the downstream CDN

More precisely, in [I-D.ietf-cdni-framework] the request routing interface is broadly divided into two functionalities:

- o 1) the asynchronous advertisement of footprint and capabilities by a dCDN that allows a uCDN to decide whether to redirect particular user requests to that dCDN (the CDNI FCI)
- o 2) the synchronous operation of actually redirecting a user request (the CDNI RI)

Application Layer Traffic Optimization (ALTO) is an approach for guiding the resource provider selection process in distributed applications that can choose among several candidate resources providers to retrieve a given resource. By conveying network layer (topology) information, an ALTO server can provide important information to "guide" the resource provider selection process in distributed applications. Usually, it is assumed that an ALTO server conveys information these applications cannot measure themselves [RFC5693].

Originally, ALTO was motivated by the huge amount of cross-ISP traffic generated by P2P applications [RFC5693]. Recently, however, ALTO is also being considered for improving the request routing in CDNs [I-D.jenkins-alto-cdn-use-cases]. In this context, it has also been proposed to use ALTO for selecting an entry-point in a downstream NSP's network (see section 3.4 "CDN delivering Over-The-Top of a NSP's network" in [I-D.jenkins-alto-cdn-use-cases]). Also, the CDNI problem statement explicitly mentions ALTO as a candidate protocol for "algorithms for selection of CDN or Surrogate by Request-Routing systems" [I-D.ietf-cdni-problem-statement]. Yet, there have not been concrete proposals so far on how to use ALTO in the context of CDN interconnection. This document tries to close this gap by giving some examples on how ALTO could be used within CDNI request routing.

3. Assumptions and High-Level Design Considerations

In this section we list some assumptions and design issues to be considered when using ALTO for the CDNI Footprint and Capabilities Advertisement interface

3.1. General Assumptions and Consideration

Below we list some general assumptions and considerations:

- o As explicitly being out-of-scope for CDNI [I-D.ietf-cdni-problem-statement], the examples used in this document assume that ingestion of content or acquiring content

across CDNs is not part of request routing as considered within CDNI standardization work. The focus of using ALTO (as considered in this document) is hence on request routing only, assuming that the content (desired by the end user) is available in the downstream CDN (or can be acquired by the downstream CDN by some means).

- o Federation Model: "Footprint and Capabilities Advertisement" and in general CDN request routing depends on the federation model among the CDN providers. Designing a suitable solution thus depends on whether a solution is needed for different settings, where CDNs consist of both NSP CDNs (serving individual ASes) and general, traditional CDNs (such as Akamai). We assume that CDNI is not designed for a setting where only NSP CDNs each serve a single AS only.
- o In this document, we assume that the upstream CDN (uCDN) makes the decision on selecting a downstream CDN, based on information that each downstream CDN has made available to the upstream CDN. Further, we assume that in principle more than one dCDN may be suitable for a given end-user request (i.e. different dCDNs may claim "overlapping" footprints). The uCDN hence potentially has to select among several candidate downstream CDNs for a given end user request.
- o It is not clear what kind(s) of business, contract, and operational relationships two peering CDNs may form. For the Internet, we see provider-customer and peering as two main relations; providers may use different charging models (e.g., 95-percentile, total volume) and may provide different SLAs. Given such unknown characteristics of CDN peering business agreements, we should design the protocol to support as much diverse potential business and operational models as possible.

3.2. Semantics for Footprint/Capabilities Advertisement

The CDNI document on "Footprint and Capabilities Semantics" [I-D.spp-cdni-rr-foot-cap-semantics] defines the semantics for the CDNI FCI. It thus provides guidance on what Footprint and Capabilities mean in a CDNI context and how a protocol solution should in principle look like. Here we briefly summarize the key points of the semantics of Footprint and Capabilities (for a detailed discussion, the reader is referred to [I-D.spp-cdni-rr-foot-cap-semantics]):

- o Often, footprint and capabilities are tied together and cannot be interpreted independently from each other. In such cases, i.e. where capabilities must be expressed on a per footprint basis, it

may be beneficial to combine footprint and capabilities advertisement.

- o Given that a large part of Footprint and Capabilities Advertisement will actually happen in contractual agreements, the semantics of CDNI Footprint and Capabilities advertisement refer to answering the following question: what exactly still needs to be advertised by the CDNI FCI? For instance, updates about temporal failures of part of a footprint can be useful information to convey via the CDNI request routing interface. Such information would provide updates on information previously agreed in contracts between the participating CDNs. In other words, the CDNI FCI is a means for a dCDN to provide changes/updates regarding a footprint and/or capabilities it has prior agreed to serve in a contract with a uCDN.
- o It seems clear that "coverage/reachability" types of footprint must be supported within CDNI. The following such types of footprint are mandatory and must be supported by the CDNI FCI:

- * List of ISO Country Codes

- * List of AS numbers

- * Set of IP-prefixes

A 'set of IP-prefixes' must be able to contain full IP addresses, i.e., a /32 for IPv4 and a /128 for IPv6, and also IP prefixes with an arbitrary prefix length. There must also be support for multiple IP address versions, i.e., IPv4 and IPv6, in such a footprint.

- o For all of these mandatory-to-implement footprint types, footprints can be viewed as constraints for delegating requests to a dCDN: A dCDN footprint advertisement tells the uCDN the limitations for delegating a request to the dCDN. For IP prefixes or ASN(s), the footprint signals to the uCDN that it should consider the dCDN a candidate only if the IP address of the request routing source falls within the prefix set (or ASN, respectively). The CDNI specifications do not define how a given uCDN determines what address ranges are in a particular ASN. Similarly, for country codes a uCDN should only consider the dCDN a candidate if it covers the country of the request routing source. The CDNI specifications do not define how a given uCDN determines the country of the request routing source. Multiple footprint constraints are additive, i.e. the advertisement of different types of footprint narrows the dCDN candidacy cumulatively.

- o The following capabilities seem useful as 'base' capabilities, i.e. ones that are needed in any case and therefore constitute mandatory capabilities to be supported by the CDNI FCI:
 - * Delivery Protocol (e.g., HTTP vs. RTMP)
 - * Acquisition Protocol (for acquiring content from a uCDN)
 - * Redirection Mode (e.g., DNS Redirection vs. HTTP Redirection as discussed in [I-D.ietf-cdni-framework])
 - * Capabilities related to CDNI Logging (e.g., supported logging mechanisms)
 - * Capabilities related to CDNI Metadata (e.g., authorization algorithms or support for proprietary vendor metadata)

4. Selection of a Downstream CDN with ALTO

Under the considerations stated in Section 3, ALTO can help the upstream CDN provider to select a proper downstream CDN provider for a given end user request as follows: Each downstream CDN provider hosts an ALTO server which provides ALTO information (i.e. ALTO network maps and ALTO cost maps [I-D.ietf-alto-protocol]) to an ALTO client at the upstream CDN provider. Network maps provided by each of several candidate downstream CDNs can provide information to the upstream CDN provider about each dCDN's "coverage/reachability" as well as capabilities.

4.1. Footprint and Capabilities Advertisement using ALTO Network Map and PID Properties

Conceptually, the footprint and capabilities interface of a dCDN is easy to specify: It is a function that given an endhost, returns if the dCDN is willing to serve the endhost, and the capabilities available to that endhost (e.g., "delivery-protocol": ["HTTP", "RTMP"], "acquisition-protocol": ["HTTP"], "redirection-mode": ["HTTP-redirect"], "logging-mechanism": ["TBD"], and "meta-capabilities": []).

Specifying the preceding for each endhost can be redundant, and one may use PIDs defined in ALTO. Specifically, an ALTO network map contains a "set of Network Location groupings" [I-D.ietf-alto-protocol]. The groupings are defined in the form of so-called "PIDs". A PID is an identifier to group network location endpoints, e.g. IP-addresses in the form of prefixes (see section 4 in [I-D.ietf-alto-protocol] for details).

Applying the basic idea of ALTO PIDs to the preceding, abstract mapping specification, by aggregating endhosts with the same capabilities in the same PID, we obtain CDNI FCI using ALTO Network Maps as simply (1) a Network Map which defines a set of PIDs, and (2) a PID Property Map [draft-roome-alto-pid-properties] that defines the properties of each PID, where the properties define the capabilities.

With the preceding Network Map and PID Property Map, the upstream CDN provider can easily match a given end user request with the footprint and capabilities of the downstream CDN providers. Whenever the footprint and/or capabilities of a dCDN change, the ALTO server of the dCDN changes its data, and the uCDN can obtain the update through ALTO incremental updates. Future extensions to ALTO to add notifications can be integrated when they become available.

In particular, this document does not define how a dCDN aggregates the endhosts into PIDs, to allow flexibility in (anticipated) updates.

In this document, we define the following PID properties, which each must be a JSON array, to convey all mandatory capabilities (see Section 3.2):

- o delivery-protocol
- o acquisition-protocol
- o redirection-mode
- o login-mechanism
- o meta-capabilities

To complement the preceding capabilities mapping, we require that an uCDN has access to ALTO Network Map(s) that can map from an endhost to Country Code and AS Number. Such mapping may or may not be specific to CDNI but can be a general mapping. Specifically, the uCDN should have access to ALTO Network Map(s) with Properties include:

- o country-code
- o asn

4.2. Conveying additional information with ALTO Cost Maps

An ALTO cost map contains costs between defined groupings of a corresponding network map (i.e. costs between PIDs): "An ALTO Cost Map defines Path Costs pairwise amongst sets of source and destination Network Locations" [I-D.ietf-alto-protocol]. This concept enables the provider of a cost map to express (and quantify) preferences of a destination network location with respect to a given source network location.

In the context of CDNI, the ALTO cost map concept is an extensive tool to convey additional information about the footprint or capabilities of a downstream CDN. The cost map concept provides a means for a downstream CDN provider to convey numeric values associated with a PID, e.g. in order to convey metrics associated with a footprint or a capability. This may be useful for future, non-mandatory types of footprint or capabilities.

One way to use ALTO cost maps would have these maps of the type N-to-m, i.e. 'costs' are expressed for each of N end user source PIDs to m dCDN request router PIDs. Semantically, a source PID in a CDNI ALTO cost map is thus the end user location, whereas a destination PID is a (group of) request router(s) to which the uCDN redirects the end user request. Note that this perspective is driven by the CDNI request routing. An alternative way - seen from the perspective of content retrieval - would be to have a m-to-N cost map where the source is always the dCDN and the destination is the end user (with the semantic "if the source dCDN would deliver content to an end user in the destination PID, the costs would be the following). With explicit destination PIDs reflecting different entries to the same dCDN, the dCDN can convey shortcut or differentiated quality of services.

4.3. Example of Selecting a Downstream CDN based on ALTO Maps

In the following, we will outline an example of dCDN selection by a uCDN based on ALTO maps provided by dCDNs. Consider the following example: An upstream CDN (uCDN) has agreed on CDN interconnection with several downstream CDNs (dCDN-a, dCDN-b, and dCDN-c). Each of these downstream CDNs runs an ALTO server to provide aforementioned ALTO information. Whenever the upstream CDN receives a request from an end user and has determined that this request is best served by an interconnected dCDN, the uCDN uses ALTO maps to make a redirection decision. For a given request, assume that only the ALTO network maps provided by dCDN-a and dCDN-c include the endhost. The uCDN first looks up the PIDs of the endhost in the two network maps from the two dCDNs, then search the PID properties to find out the capabilities of each dCDN for the endhost. If only one dCDN supports the required capabilities, then the uCDN chooses the dCDN. Otherwise, if Cost Maps are available to provide additional server

selection information (e.g., a Cost Map defining latency), the uCDN picks the dCDN with better cost performance.

4.4. Advantages of using ALTO

The following reasons make ALTO a suitable candidate protocol for downstream CDN selection as part of CDNI request routing and in particular for a FCI protocol:

- o CDN request routing is done at the application layer. ALTO is a protocol specifically designed to improve application layer traffic (and application layer connections among hosts on the Internet) by providing additional information to applications that these applications could not easily retrieve themselves. For CDNI, this is exactly the case: a uCDN wants to improve application layer CDN request routing by using dedicated information (provided by a dCDN) that the uCDN could not easily obtain otherwise.
- o The semantics of an ALTO network are an exact match for the needed information to convey a footprint by a downstream CDN, in particular if such a footprint is being expressed by IP-prefix ranges.
- o ALTO cost maps are suitable to express various types of numeric values and can hence be used by an upstream CDN to obtain metrics for capabilities associated with a given dCDN for a given footprint. Further, an ALTO cost map could also convey relevant network topology information other than simply routing hops or reachability. This facilitates advanced and more sophisticated selection of a downstream CDN based on various metrics by the upstream CDN and increases flexibility to cover different use cases and business models for CDN interconnection.
- o Flexible granularity: The concept of the PID and ALTO network/cost maps allows for different degrees of granularity. This enables a dCDN to differentiate the delivery quality for serving an end user request on a fine granularity depending on the end user location (and not only express delivery quality e.g. on an AS-level). It remains at the discretion of each dCDN how fine-granular the ALTO network and cost maps are that it publishes.
- o ALTO maps can be signed and hence provide inherent integrity protection (see Section 6)

5. Useful ALTO extensions for CDNI Request Routing

It is envisioned that yet-to-be-defined ALTO extensions will be standardized that make the ALTO protocol more suitable and useful for applications other than the originally considered P2P use case [I-D.marocco-alto-next]. Some of these extensions to the ALTO protocol would be useful for ALTO to be used as a protocol within CDNI request routing, and in particular within the "Footprint and Capabilities Advertisement" part of the CDNI request routing interface.

The following proposed extensions to ALTO would be beneficial to facilitate CDNI request routing with ALTO as outlined in Section 4:

- o Server-initiated Notifications and Incremental Updates: In case the footprint or the capabilities of a downstream CDN change abruptly (i.e. unexpectedly from the perspective of an upstream CDN), server initiated notifications would enable a dCDN to directly inform an upstream CDN about such changes. Consider the case where - due to failure - part of the footprint of the dCDN is not functioning, i.e. the CDN cannot serve content to such clients with reasonable QoS. Without server-initiated notifications, the uCDN might still use a very recent network and cost map from dCDN, and therefore redirect request to dCDN which it cannot serve. Similarly, the possibility for incremental updates would enable efficient conveyance of the aforementioned (or similar) status changes by the dCDN to the uCDN. A proposal for server-initiated ALTO updates can be found in [I-D.marocco-alto-ws]. A discussion of incremental ALTO updates can be found in [I-D.schwan-alto-incr-updates].
- o Content Availability on Hosts: A dCDN might want to express CDN capabilities in terms of certain content types (e.g. codecs/formats, or content from certain content providers). A new endpoint property for ALTO that would be able to express such "content availability" would enable a dCDN to make available such information to an upstream CDN. This would enable a uCDN to determine if a given dCDN actually has the capabilities for a given request with respect to the type of content requested.

- o Resource Availability on Hosts or Links: The capabilities on links (e.g. maximum bandwidth) or caches (e.g. average load) might be useful information for an upstream CDN for optimized downstream CDN selection. For instance, if a uCDN receives a streaming request for content with a certain bitrate, it needs to know if it is likely that a dCDN can fulfill such stringent application-level requirements (i.e. can be expected to have enough consistent bandwidth) before it redirects the request. In general, if ALTO could convey such information via new endpoint properties, it would enable more sophisticated means for downstream CDN selection with ALTO.

6. Security Considerations

One important security consideration is the proper authentication of advertisement information provided by a downstream CDN. The ALTO protocol provides a specification for a signature of ALTO maps (see 8.2.2. of [I-D.ietf-alto-protocol]). ALTO thus provides a proper means for protecting the integrity of footprint advertisement information.

More Security Considerations will be discussed in a future version of this document.

7. Summary and Outlook

This document presented concrete examples of how ALTO can be used within the downstream CDN selection of CDNI Request Routing. Further, the document provides arguments why ALTO is a meaningful protocol in this context. Essentially, ALTO network and cost maps are a means to provide detailed and various types of information to an upstream CDN, in order to facilitate well-considered downstream CDN selection.

The intention of this document is to find consensus in the CDNI WG that ALTO is a useful protocol for CDNI request routing, and that ALTO has many benefits for proper selection of a downstream CDN. The overall objective is to form agreement on how ALTO should be used within the CDNI request routing protocol. It is the intention to capture the outcome of such continuing discussions in future versions of this document.

8. Acknowledgements

Jan Seedorf is partially supported by the CHANGE project (CHANGE: Enabling Innovation in the Internet Architecture through Flexible Flow-Processing Extensions, <http://www.change-project.eu/>), a research project supported by the European Commission under its 7th

Framework Program (contract no. 257422). The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the CHANGE project or the European Commission.

Jan Seedorf has been partially supported by the COAST project (Content Aware Searching, retrieval and sTreaming, <http://www.coast-fp7.eu>), a research project supported by the European Commission under its 7th Framework Program (contract no. 248036). The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the COAST project or the European Commission.

9. References

9.1. Normative References

- [RFC5693] Seedorf, J. and E. Burger, "Application-Layer Traffic Optimization (ALTO) Problem Statement", RFC 5693, October 2009.
- [RFC6707] Niven-Jenkins, B., Le Faucheur, F., and N. Bitar, "Content Distribution Network Interconnection (CDNI) Problem Statement", RFC 6707, September 2012.
- [RFC6770] Bertrand, G., Stephan, E., Burbridge, T., Eardley, P., Ma, K., and G. Watson, "Use Cases for Content Delivery Network Interconnection", RFC 6770, November 2012.

9.2. Informative References

- [I-D.peterson-cdni-strawman]
Peterson, L. and J. Hartman, "Content Distribution Network Interconnection (CDNI) Problem Statement", draft-peterson-cdni-strawman-01 (work in progress), May 2011.
- [I-D.ietf-cdni-problem-statement]
Niven-Jenkins, B., Faucheur, F., and N. Bitar, "Content Distribution Network Interconnection (CDNI) Problem Statement", draft-ietf-cdni-problem-statement-08 (work in progress), June 2012.
- [I-D.marocco-alto-next]
Marocco, E. and V. Gurbani, "Extending the Application-Layer Traffic Optimization (ALTO) Protocol", draft-marocco-alto-next-00 (work in progress), January 2012.

- [I-D.ietf-alto-protocol]
Alimi, R., Penno, R., and Y. Yang, "ALTO Protocol", draft-ietf-alto-protocol-20 (work in progress), October 2013.
- [I-D.ietf-cdni-requirements]
Leung, K. and Y. Lee, "Content Distribution Network Interconnection (CDNI) Requirements", draft-ietf-cdni-requirements-11 (work in progress), October 2013.
- [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-10 (work in progress), August 2012.
- [I-D.marocco-alto-ws]
Marocco, E. and J. Seedorf, "WebSocket-based server-to-client notifications for the Application-Layer Traffic Optimization (ALTO) Protocol", draft-marocco-alto-ws-01 (work in progress), July 2012.
- [I-D.schwan-alto-incr-updates]
Schwan, N. and B. Roome, "ALTO Incremental Updates", draft-schwan-alto-incr-updates-02 (work in progress), July 2012.
- [I-D.jenkins-alto-cdn-use-cases]
Niven-Jenkins, B., Watson, G., Bitar, N., Medved, J., and S. Previdi, "Use Cases for ALTO within CDNs", draft-jenkins-alto-cdn-use-cases-03 (work in progress), June 2012.
- [I-D.seedorf-alto-for-cdni]
Seedorf, J., "ALTO for CDNI Request Routing", draft-seedorf-alto-for-cdni-00 (work in progress), October 2011.
- [I-D.ietf-cdni-framework]
Peterson, L. and B. Davie, "Framework for CDN Interconnection", draft-ietf-cdni-framework-06 (work in progress), October 2013.
- [I-D.liu-cdni-cost]
Liu, H., "A Cost Perspective on Using Multiple CDNs", draft-liu-cdni-cost-00 (work in progress), October 2011.
- [I-D.spp-cdni-rr-foot-cap-semantics]
Seedorf, J., Peterson, J., Previdi, S., Brandenburg, R., and K. Ma, "CDNI Request Routing: Footprint and

Capabilities Semantics", draft-spp-cdni-rr-foot-cap-
semantics-04 (work in progress), February 2013.

Authors' Addresses

Jan Seedorf
NEC Laboratories Europe, NEC Europe Ltd.
Kurfuersten-Anlage 36
Heidelberg 69115
Germany

Phone: +49 (0) 6221 4342 221
Email: jan.seedorf@neclab.eu
URI: <http://www.neclab.eu>

Y.R. Yang
Yale University
51 Prospect Street
New Haven 06511
USA

Email: yry@cs.yale.edu
URI: <http://www.cs.yale.edu/~yry/>

CDNI	J. Seedorf
Internet-Draft	HFT Stuttgart - Univ. of Applied Sciences
Intended status: Standards Track	Y. Yang
Expires: January 3, 2018	Tongji/Yale
	K. Ma
	Ericsson
	J. Peterson
	Neustar
	July 2, 2017

Content Delivery Network Interconnection (CDNI) Request Routing: CDNI
Footprint and Capabilities Advertisement using ALTO
draft-seedorf-cdni-request-routing-alto-10

Abstract

The Content Delivery Networks Interconnection (CDNI) WG is defining a set of protocols to inter-connect CDNs, to achieve multiple goals such as extending the reach of a given CDN to areas that are not covered by that particular CDN. One component that is needed to achieve the goal of CDNI is the CDNI Request Routing Footprint & Capabilities Advertisement interface (FCI) [RFC7336]. [RFC8008] has defined precisely the semantics of FCI and provided guidelines on the FCI protocol, but the exact protocol is explicitly outside the scope of that document. In this document, we define an FCI protocol using the Application Layer Traffic Optimization (ALTO) protocol.

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 3, 2018.

Copyright Notice

Copyright (c) 2017 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	2
2. Background	4
2.1. Semantics of FCI Advertisement	4
2.2. ALTO Background and Benefits	5
3. CDNI FCI ALTO Service	7
3.1. Server Response Encoding	8
3.1.1. Media Type	8
3.1.2. Meta Information	8
3.1.3. Data Information	8
3.2. Protocol Errors	8
3.3. Examples	8
3.3.1. Basic Example	8
3.3.2. Incremental FCI Update Example	9
3.3.3. FCI Using ALTO Network Map Example	9
4. Security Considerations	9
5. Acknowledgements	10
6. References	10
6.1. Normative References	10
6.2. Informative References	11
Authors' Addresses	11

1. Introduction

Many Network Service Providers (NSPs) are currently considering or have already started to deploy Content Delivery Networks (CDNs) within their networks. As a consequence of this development, there is a need for interconnecting these local CDNs. Content Delivery Networks Interconnection (CDNI) has the goal of standardizing protocols to enable such interconnection of CDNs [RFC6707].

The CDNI problem statement [RFC6707] defines four interfaces to be standardized within the IETF for CDN interconnection:

- o CDNI Request Routing Interface
- o CDNI Metadata Interface
- o CDNI Logging Interface
- o CDNI Control Interface

The main purpose of the CDNI Request Routing Interface is described in [RFC6707] as follows: "The CDNI Request Routing interface enables a Request Routing function in an Upstream CDN to query a Request Routing function in a Downstream CDN to determine if the Downstream CDN is able (and willing) to accept the delegated Content Request. It also allows the Downstream CDN to control what should be returned to the User Agent in the redirection message by the upstream Request Routing function." On a high level, the scope of the CDNI Request Routing Interface therefore contains two main tasks:

- o determining if the downstream CDN is willing to accept a delegated content request;
- o redirecting the content request coming from an upstream CDN to the proper entry point or entity in the downstream CDN.

Correspondingly, the request routing interface is broadly divided into two functionalities:

- o CDNI FCI: the advertisement from a dCDN to a uCDN or a query from a uCDN to a dCDN for the uCDN to decide whether to redirect particular user requests to that dCDN;
- o CDNI RI: the synchronous operation of actually redirecting a user request.

This document focuses solely on CDNI FCI, with a goal to specify a new Application Layer Traffic Optimization (ALTO) [RFC7285] service called 'CDNI/FCI Service', to transport and update CDNI FCI JSON objects, which are defined in a separate document in [RFC8008].

Throughout this document, we use the terminology for CDNI defined in [RFC6707] and [RFC8008].

2. Background

The design of CDNI FCI transport using ALTO depends on understanding of both FCI semantics and ALTO. Hence, we start with a review of both.

2.1. Semantics of FCI Advertisement

The CDNI document on "Footprint and Capabilities Semantics" [RFC8008] defines the semantics for the CDNI FCI. It thus provides guidance on what Footprint and Capabilities mean in a CDNI context and how a protocol solution should in principle look like. The definitions in [RFC8008] depend on [RFC8006]. Here we briefly summarize key related points of [RFC8008] and [RFC8006]. For a detailed discussion, the reader is referred to the RFCs.

- o Footprint and capabilities are tied together and cannot be interpreted independently from each other. In such cases, i.e. where capabilities must be expressed on a per footprint basis, it may be beneficial to combine footprint and capabilities advertisement. [RFC8008] integrates footprint and capabilities with an approach of "capabilities with footprint restrictions".
- o Given that a large part of Footprint and Capabilities Advertisement will actually happen in contractual agreements, the semantics of CDNI Footprint and Capabilities advertisement refer to answering the following question: what exactly still needs to be advertised by the CDNI FCI? For instance, updates about temporal failures of part of a footprint can be useful information to convey via the CDNI request routing interface. Such information would provide updates on information previously agreed in contracts between the participating CDNs. In other words, the CDNI FCI is a means for a dCDN to provide changes/updates regarding a footprint and/or capabilities it has prior agreed to serve in a contract with a uCDN. Hence, server push and incremental encoding will be necessary techniques.
- o Multiple types of footprints are defined in [RFC8006]:
 - * List of ISO Country Codes
 - * List of AS numbers
 - * Set of IP-prefixes

A 'set of IP-prefixes' must be able to contain full IP addresses, i.e., a /32 for IPv4 and a /128 for IPv6, and also IP prefixes with an arbitrary prefix length. There must also be support for

multiple IP address versions, i.e., IPv4 and IPv6, in such a footprint.

- o For all of these mandatory-to-implement footprint types, footprints can be viewed as constraints for delegating requests to a dCDN: A dCDN footprint advertisement tells the uCDN the limitations for delegating a request to the dCDN. For IP prefixes or ASN(s), the footprint signals to the uCDN that it should consider the dCDN a candidate only if the IP address of the request routing source falls within the prefix set (or ASN, respectively). The CDNI specifications do not define how a given uCDN determines what address ranges are in a particular ASN. Similarly, for country codes a uCDN should only consider the dCDN a candidate if it covers the country of the request routing source. The CDNI specifications do not define how a given uCDN determines the country of the request routing source. Multiple footprint constraints are additive, i.e. the advertisement of different types of footprint narrows the dCDN candidacy cumulatively.
- o The following capabilities are defined as 'base' capabilities, i.e. ones that are needed in any case and therefore constitute mandatory capabilities to be supported by the CDNI FCI:
 - * Delivery Protocol (e.g., HTTP vs. RTMP)
 - * Acquisition Protocol (for acquiring content from a uCDN)
 - * Redirection Mode (e.g., DNS Redirection vs. HTTP Redirection as discussed in [RFC7336])
 - * Capabilities related to CDNI Logging (e.g., supported logging mechanisms)
 - * Capabilities related to CDNI Metadata (e.g., authorization algorithms or support for proprietary vendor metadata)

2.2. ALTO Background and Benefits

Application Layer Traffic Optimization (ALTO) [RFC7285] is an approach for guiding the resource provider selection process in distributed applications that can choose among several candidate resources providers to retrieve a given resource. By conveying network layer (topology) information, an ALTO server can provide important information to "guide" the resource provider selection process in distributed applications. Usually, it is assumed that an ALTO server conveys information these applications cannot measure themselves [RFC5693].

Originally, ALTO was motivated by the huge amount of cross-ISP traffic generated by P2P applications [RFC5693]. Recently, however, ALTO is also being considered for improving the request routing in CDNs [I-D.jenkins-alto-cdn-use-cases]. In this context, it has also been proposed to use ALTO for selecting an entry-point in a downstream NSP's network (see section 3.4 "CDN delivering Over-The-Top of a NSP's network" in [I-D.jenkins-alto-cdn-use-cases]). Also, the CDNI problem statement explicitly mentions ALTO as a candidate protocol for "algorithms for selection of CDN or Surrogate by Request-Routing systems" [RFC6707].

The following reasons make ALTO a suitable candidate protocol for downstream CDN selection as part of CDNI request routing and in particular for an FCI protocol:

- o CDN request routing is done at the application layer. ALTO is a protocol specifically designed to improve application layer traffic (and application layer connections among hosts on the Internet) by providing additional information to applications that these applications could not easily retrieve themselves. For CDNI, this is exactly the case: a uCDN wants to improve application layer CDN request routing by using dedicated information (provided by a dCDN) that the uCDN could not easily obtain otherwise.
- o The semantics of an ALTO network map are an exact match for the needed information to convey a footprint by a downstream CDN, in particular if such a footprint is being expressed by IP-prefix ranges.
- o Security: ALTO maps can be signed and hence provide inherent integrity protection (see Section 4)
- o RESTful-Design: The ALTO protocol has undergone extensive revisions in order to provide a RESTful design regarding the client-server interaction specified by the protocol. A CDNI FCI interface based on ALTO would inherit this RESTful design.
- o Error-handling: The ALTO protocol has undergone extensive revisions in order to provide sophisticated error-handling, in particular regarding unexpected cases. A CDNI FCI interface based on ALTO would inherit this thought-through and mature error-handling.
- o Filtered network map: The ALTO Map Filtering Service (see [RFC7285] for details) would allow a uCDN to query only for parts of an ALTO map.

- o Server-initiated Notifications and Incremental Updates: In case the footprint or the capabilities of a downstream CDN change abruptly (i.e. unexpectedly from the perspective of an upstream CDN), server initiated notifications would enable a dCDN to directly inform an upstream CDN about such changes. Consider the case where - due to failure - part of the footprint of the dCDN is not functioning, i.e. the CDN cannot serve content to such clients with reasonable QoS. Without server-initiated notifications, the uCDN might still use a very recent network and cost map from dCDN, and therefore redirect request to dCDN which it cannot serve. Similarly, the possibility for incremental updates would enable efficient conveyance of the aforementioned (or similar) status changes by the dCDN to the uCDN. The newest design of ALTO supports server pushed incremental updates [I-D.ietf-alto-incr-update-sse].
- o Content Availability on Hosts: A dCDN might want to express CDN capabilities in terms of certain content types (e.g. codecs/formats, or content from certain content providers). The new endpoint property for ALTO would enable a dCDN to make available such information to an upstream CDN. This would enable a uCDN to determine if a given dCDN actually has the capabilities for a given request with respect to the type of content requested.
- o Resource Availability on Hosts or Links: The capabilities on links (e.g. maximum bandwidth) or caches (e.g. average load) might be useful information for an upstream CDN for optimized downstream CDN selection. For instance, if a uCDN receives a streaming request for content with a certain bitrate, it needs to know if it is likely that a dCDN can fulfill such stringent application-level requirements (i.e. can be expected to have enough consistent bandwidth) before it redirects the request. In general, if ALTO could convey such information via new endpoint properties, it would enable more sophisticated means for downstream CDN selection with ALTO.

3. CDNI FCI ALTO Service

The ALTO protocol is based on an ALTO Information Service Framework which consists of several services, where all ALTO services are 'provided through a common transport protocol, messaging structure and encoding, and transaction model' [RFC7285]. The ALTO protocol specification [RFC7285] defines several such services, e.g. the ALTO map service.

This document defines a new ALTO Service called 'CDNI Footprint & Capabilities Advertisement Service' which conveys JSON objects of media type 'application/cdni'. This media type and JSON object

format is defined in [RFC8006] and [RFC8008]; this document specifies how to transport such JSON objects via the ALTO protocol with the ALTO 'CDNI Footprint & Capabilities Advertisement Service'.

3.1. Server Response Encoding

3.1.1. Media Type

The media type of the CDNI FCI Map is 'application/cdni'.

3.1.2. Meta Information

The 'meta' field of a FCI response MUST include 'vtag', which is an ALTO Version Tag of the retrieved FCIMapData according to [RFC7285] (Section 10.3.). It thus contains a 'resource-id' attribute, and a 'tag' is an identifier string.

3.1.3. Data Information

The data component of a CDNI FCI resource is named 'cdni-fcimap' which is a JSON object defined by [RFC8008]. This JSON object is derived from ResponseEntityBase as specified in the ALTO protocol [RFC7285] (Section 8.4.).

3.2. Protocol Errors

Protocol errors are handled as specified in the ALTO protocol [RFC7285] (Section 8.5.).

3.3. Examples

3.3.1. Basic Example

The following example shows an CDNI FCI response as in [RFC8008], however with meta-information as defined in Section 3.1.2 of this document.

```
GET /fcimap HTTP/1.1
Host: alto.example.com
Accept: application/cdni,application/alto-error+json

HTTP/1.1 200 OK
Content-Length: 439
Content-Type: application/cdni
{
  "meta" : {
    "vtag": {
      "resource-id": "my-default-fcimap",
      "tag": "da65eca2eb7a10ce8b059740b0b2e3f8eb1d4785"
    }
  },
  "cdni-fcimap": {
    "capabilities": [
      {
        "capability-type": "FCI.DeliveryProtocol",
        "capability-value": {
          "delivery-protocols": [
            "http/1.1",
          ]
        },
        "footprints": [
          <Footprint objects>
        ]
      }
    ]
  }
}
```

3.3.2. Incremental FCI Update Example

3.3.3. FCI Using ALTO Network Map Example

4. Security Considerations

One important security consideration is the proper authentication of advertisement information provided by a downstream CDN. The ALTO protocol provides a specification for a signature of ALTO information (see 8.2.2. of [RFC7285]). ALTO thus provides a proper means for protecting the integrity of FCI information.

More Security Considerations will be discussed in a future version of this document.

5. Acknowledgements

The authors would like to thank Kevin Ma, Daryl Malas, and Matt Caulfield for their timely reviews and invaluable comments.

Jan Seedorf is partially supported by the GreenICN project (GreenICN: Architecture and Applications of Green Information Centric Networking), a research project supported jointly by the European Commission under its 7th Framework Program (contract no. 608518) and the National Institute of Information and Communications Technology (NICT) in Japan (contract no. 167). The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the GreenICN project, the European Commission, or NICT.

6. References

6.1. Normative References

- [RFC5693] Seedorf, J. and E. Burger, "Application-Layer Traffic Optimization (ALTO) Problem Statement", RFC 5693, DOI 10.17487/RFC5693, October 2009, <<http://www.rfc-editor.org/info/rfc5693>>.
- [RFC6707] Niven-Jenkins, B., Le Faucheur, F., and N. Bitar, "Content Distribution Network Interconnection (CDNI) Problem Statement", RFC 6707, DOI 10.17487/RFC6707, September 2012, <<http://www.rfc-editor.org/info/rfc6707>>.
- [RFC6770] Bertrand, G., Ed., Stephan, E., Burbidge, T., Eardley, P., Ma, K., and G. Watson, "Use Cases for Content Delivery Network Interconnection", RFC 6770, DOI 10.17487/RFC6770, November 2012, <<http://www.rfc-editor.org/info/rfc6770>>.
- [RFC7285] Alimi, R., Ed., Penno, R., Ed., Yang, Y., Ed., Kiesel, S., Previdi, S., Roome, W., Shalunov, S., and R. Woundy, "Application-Layer Traffic Optimization (ALTO) Protocol", RFC 7285, DOI 10.17487/RFC7285, September 2014, <<http://www.rfc-editor.org/info/rfc7285>>.
- [RFC7336] Peterson, L., Davie, B., and R. van Brandenburg, Ed., "Framework for Content Distribution Network Interconnection (CDNI)", RFC 7336, DOI 10.17487/RFC7336, August 2014, <<http://www.rfc-editor.org/info/rfc7336>>.

- [RFC7337] Leung, K., Ed. and Y. Lee, Ed., "Content Distribution Network Interconnection (CDNI) Requirements", RFC 7337, DOI 10.17487/RFC7337, August 2014, <<http://www.rfc-editor.org/info/rfc7337>>.
- [RFC8006] Niven-Jenkins, B., Murray, R., Caulfield, M., and K. Ma, "Content Delivery Network Interconnection (CDNI) Metadata", RFC 8006, DOI 10.17487/RFC8006, December 2016, <<http://www.rfc-editor.org/info/rfc8006>>.
- [RFC8008] Seedorf, J., Peterson, J., Previdi, S., van Brandenburg, R., and K. Ma, "Content Delivery Network Interconnection (CDNI) Request Routing: Footprint and Capabilities Semantics", RFC 8008, DOI 10.17487/RFC8008, December 2016, <<http://www.rfc-editor.org/info/rfc8008>>.

6.2. Informative References

- [I-D.ietf-alto-incr-update-sse]
Roome, W. and Y. Yang, "ALTO Incremental Updates Using Server-Sent Events (SSE)", draft-ietf-alto-incr-update-sse-07 (work in progress), July 2017.
- [I-D.jenkins-alto-cdn-use-cases]
Niven-Jenkins, B., Watson, G., Bitar, N., Medved, J., and S. Previdi, "Use Cases for ALTO within CDNs", draft-jenkins-alto-cdn-use-cases-03 (work in progress), June 2012.
- [I-D.ma-cdni-capabilities]
Ma, K. and J. Seedorf, "CDNI Footprint & Capabilities Advertisement Interface", draft-ma-cdni-capabilities-09 (work in progress), April 2016.

Authors' Addresses

Jan Seedorf
HFT Stuttgart - Univ. of Applied Sciences
Schellingstrasse 24
Stuttgart 70174
Germany

Phone: +49-0711-8926-2801
Email: jan.seedorf@hft-stuttgart.de

Y.R. Yang
Tongji/Yale University
51 Prospect Street
New Haven, CT 06511
United States of America

Email: yry@cs.yale.edu
URI: <http://www.cs.yale.edu/~yry/>

Kevin J. Ma
Ericsson
43 Nagog Park
Acton, MA 01720
United States of America

Phone: +1-978-844-5100
Email: kevin.j.ma@ericsson.com

Jon Peterson
NeuStar
1800 Sutter St Suite 570
Concord, CA 94520
United States of America

Email: jon.peterson@neustar.biz