

Internet Engineering Task Force (IETF)
Request for Comments: 8008
Category: Standards Track
ISSN: 2070-1721

J. Seedorf
HFT Stuttgart - Univ. of Applied Sciences
J. Peterson
Neustar
S. Previdi
Cisco
R. van Brandenburg
TNO
K. Ma
Ericsson
December 2016

Content Delivery Network Interconnection (CDNI) Request Routing:
Footprint and Capabilities Semantics

Abstract

This document captures the semantics of the "Footprint and Capabilities Advertisement" part of the Content Delivery Network Interconnection (CDNI) Request Routing interface, i.e., the desired meaning of "Footprint" and "Capabilities" in the CDNI context and what the "Footprint & 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.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc8008>.

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	4
1.1. Terminology	5
2. Design Decisions for Footprint and Capabilities	6
2.1. Advertising Limited Coverage	6
2.2. Capabilities and Dynamic Data	7
2.3. Advertisement versus Queries	8
2.4. Avoiding or Handling "Cheating" dCDNs	8
3. Focusing on Capabilities with Footprint Restrictions	9
4. Footprint and Capabilities Extension	9
5. Capability Advertisement Object	11
5.1. Base Advertisement Object	12
5.2. Encoding	12
5.3. Delivery Protocol Capability Object	13
5.3.1. Delivery Protocol Capability Object Serialization ..	13
5.4. Acquisition Protocol Capability Object	14
5.4.1. Acquisition Protocol Capability Object Serialization	14
5.5. Redirection Mode Capability Object	15
5.5.1. Redirection Mode Capability Object Serialization ..	15
5.6. CDNI Logging Capability Object	16
5.6.1. CDNI Logging Capability Object Serialization	17
5.7. CDNI Metadata Capability Object	18
5.7.1. CDNI Metadata Capability Object Serialization	19
6. IANA Considerations	20
6.1. CDNI Payload Types	20
6.1.1. CDNI FCI DeliveryProtocol Payload Type	20
6.1.2. CDNI FCI AcquisitionProtocol Payload Type	20
6.1.3. CDNI FCI RedirectionMode Payload Type	20
6.1.4. CDNI FCI Logging Payload Type	21
6.1.5. CDNI FCI Metadata Payload Type	21
6.2. "CDNI Capabilities Redirection Modes" Registry	21
7. Security Considerations	22
8. References	23
8.1. Normative References	23
8.2. Informative References	24
Appendix A. Main Use Case to Consider	25
Appendix B. Semantics for Footprint Advertisement	25
Appendix C. Semantics for Capabilities Advertisement	27
Acknowledgments	30
Authors' Addresses	30

1. Introduction and Scope

The CDNI working group is working on a set of protocols to enable the interconnection of multiple CDNs. These CDNI protocols can serve multiple purposes, as discussed in [RFC6770] -- for instance, to extend the reach of a given CDN to areas in the network that are not covered by that 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 the FCI) [RFC7336], 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.

Deciding on specific protocols to use for the FCI is explicitly outside the scope of this document. However, we provide guidelines for such FCI protocols.

We make the following general assumptions in this document:

- o The CDNs participating in the CDN interconnection have already performed a bootstrap 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 advertisements and/or capability advertisements from a set of dCDNs. Footprint advertisements and capability advertisements need not use the same underlying protocol.
- o The uCDN receives the initial Request Routing message from the endpoint requesting the resource.

The CDNI problem statement [RFC6707] describes the Request Routing interface 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, [RFC6707] 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 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 document [RFC7337] lists the specific requirements for the CDNI FCI 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 content 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 CIDR or IPv6 CIDR (Classless Inter-Domain Routing), network ID (e.g., Autonomous System Number (ASN)), nation boundaries (e.g., country code), or GPS coordinates. This document does not define or endorse the quality or suitability of any particular footprint description method; rather, it only defines a method for transporting known footprint descriptions in Footprint and Capabilities Advertisement messages.
- o Capability: a feature of a dCDN upon whose support a uCDN relies 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; rather, it describes a scheme for defining new capabilities and how to transport them in Footprint and Capabilities Advertisement messages.

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. 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 a 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 area but seek to interconnect with other CDNs in order to create a single CDN fabric that 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 advertisements and capability advertisements 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 made for the CDNI FCI.

2.1. Advertising Limited Coverage

The basic use case that would motivate a dCDN to advertise 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 backups, 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 select 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 in knowing how optimal the coverage is from each of the dCDNs. Coverage need not be binary (i.e., 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 information 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 a footprint depend on whether (1) such qualitative metrics for expressing a footprint (such as the coverage "score" mentioned above) are included as part of the CDNI FCI or (2) the focus is just on a "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, the network where the Surrogate is deployed, the nature of the resource sought, and the administrative policies of the respective networks.

In the absence of network-layer impediments to reaching Surrogates, the choice to limit coverage is, by necessity, 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 either as 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 the footprint and the capability information 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 advertisements and capability advertisements 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 to select a certain dCDN for Request Routing based on information the uCDN has received from this particular dCDN. It can be assumed that cheating dCDNs 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.

Furthermore, understanding that in most cases contractual agreements will define the basic coverage used in delegation decisions, the primary focus of the FCI is on providing updates to the basic capabilities and coverage by the dCDNs. As such, the FCI has chosen the semantics of "capabilities with footprint restrictions".

4. Footprint and Capabilities Extension

Other optional "coverage/reachability" footprint types or "resource" footprint types 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 [RFC8006]).

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 CDNI Logging fields)
- o CDNI Metadata (i.e., supported GenericMetadata 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 that 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 [RFC7937] 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 Field Names" registry [RFC7937] defines all known CDNI Logging fields. CDNI Logging fields may be reused by different record-types and be mandatory-to-implement in some record-types, but they may be optional in other record-types. CDNs MUST advertise support for optional CDNI Logging fields within the context of a specific record-type. For a given record-type, CDNs SHOULD NOT advertise support for mandatory-to-implement CDNI Logging fields. The following CDNI Logging fields are defined as optional for the "cdni_http_request_v1" record-type [RFC7937]:

- o s-ccid
- o s-sid

[RFC8006] 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 provided. Advertising support for non-mandatory-to-enforce GenericMetadata

types SHOULD be provided. 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 footprint types and capability types implies that certain implementations might not support all kinds of footprints and capabilities. Therefore, any FCI solution protocol MUST define how the support for optional footprint types and capability types 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 footprint or capability types.

In general, a uCDN MAY ignore capabilities or footprint types it does not understand; in this case, it only selects a suitable dCDN based on the types of capabilities and footprints it understands. Similarly, if a dCDN does not use an optional capability or footprint that is, however, supported by a uCDN, this causes no problem for 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) [RFC8006] to facilitate a uniform set of mandatory parsing requirements for all future FCI objects.

Future object definitions (e.g., regarding the CDNI Metadata or CDNI Logging interfaces) 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 the capability-type property above

Mandatory-to-Specify: Yes.

Property: footprints

Description: CDNI capability footprint.

Type: List of CDNI Footprint objects (from the "CDNI Metadata Footprint Types" registry [RFC8006])

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 (Internet 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 [RFC8006]).

Property: delivery-protocols

Description: List of supported CDNI delivery protocols.

Type: List of protocol types (from the "CDNI Metadata Protocol Types" registry [RFC8006])

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 Transport Layer Security (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 [RFC8006]).

Property: acquisition-protocols

Description: List of supported CDNI acquisition protocols.

Type: List of protocol types (from the "CDNI Metadata Protocol Types" registry [RFC8006])

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 the "CDNI Capabilities Redirection Modes" registry, defined in 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 (i.e., not recursive) redirection with HTTP and DNS.

```
{
  "capabilities": [
    {
      "capability-type": "FCI.RedirectionMode",
      "capability-value": {
        "redirection-modes": [
          "DNS-I",
          "HTTP-I"
        ]
      }
    }
  ]
}
```

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 that are marked as optional for the specified record-types [RFC7937].

Property: record-type

Description: Supported CDNI Logging record-type.

Type: String corresponding to an entry from the "CDNI Logging record-types" registry [RFC7937]

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 [RFC7937]

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 CDNI Logging field (but not the optional Session ID CDNI 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>
      ]
    }
  ]
}
```

The final example shows the CDNI Logging capability object serialization for a CDN that supports none of the optional fields for the "cdni_http_request_v1" record-type.

```
{
  "capabilities": [
    {
      "capability-type": "FCI.Logging",
      "capability-value": {
        "record-type": "cdni_http_request_v1",
        "fields": []
      },
      "footprints": [
        <Footprint objects>
      ]
    }
  ]
}
```

5.7. CDNI Metadata Capability Object

The CDNI Metadata capability object is used to indicate support for CDNI GenericMetadata types [RFC8006].

Property: metadata

Description: List of supported CDNI GenericMetadata types.

Type: List of strings corresponding to entries from the "CDNI Payload Types" 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) [RFC8006].

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 any security policies, e.g., time, location, or protocol Access Control Lists (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 registers the following CDNI Payload Types under the IANA "CDNI Payload Types" registry:

Payload Type	Specification
FCI.DeliveryProtocol	RFC 8008
FCI.AcquisitionProtocol	RFC 8008
FCI.RedirectionMode	RFC 8008
FCI.Logging	RFC 8008
FCI.Metadata	RFC 8008

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. "CDNI Capabilities Redirection Modes" Registry

IANA has created a new "CDNI Capabilities Redirection Modes" registry in the "Content Delivery Network Interconnection (CDNI) Parameters" registry. The "CDNI Capabilities Redirection Modes" namespace defines the valid redirection modes that can be advertised as supported by a CDN. Additions to the "CDNI Capabilities Redirection Modes" 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	RFC 8008
DNS-R	Recursive DNS-based Redirection	RFC 8008
HTTP-I	Iterative HTTP-based Redirection	RFC 8008
HTTP-R	Recursive HTTP-based Redirection	RFC 8008

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 capabilities and footprint advertisement objects 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 that claims the network has no footprint or capability. This would prevent the uCDN from delegating any requests to the dCDN. Since a preexisting relationship between all dCDNs and uCDNs is assumed by CDNI, the exchange of any necessary credentials could be conducted before the FCI is brought online. The authorization decision to accept advertisements would also follow this preexisting relationship and any contractual obligations that it stipulates.

All protocols that implement these capabilities and footprint advertisement objects are REQUIRED to provide confidentiality services. Some dCDNs are willing to share information about their footprints 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 event 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

- [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>>.
- [RFC7937] Le Faucheur, F., Ed., Bertrand, G., Ed., Oprescu, I., Ed., and R. Peterkofsky, "Content Distribution Network Interconnection (CDNI) Logging Interface", RFC 7937, DOI 10.17487/RFC7937, August 2016, <<http://www.rfc-editor.org/info/rfc7937>>.
- [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>>.

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 narrow down the requirements for the CDNI FCI. To this end, the following (simplified) use case can help clarify the semantics of footprints 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 types 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, for either all of footprint 1 or only 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 a dCDN's "ability and willingness to serve". However, in addition to simple ability and willingness to serve, the uCDN could want additional information before deciding which dCDN to select, e.g., "how well" a given dCDN can actually serve a given end user request. The dCDN's ability and willingness to serve SHOULD be distinguished from the subjective qualitative measurement of how well it can serve a given end user request. One can imagine that such additional information is implicitly associated with a given footprint, due to contractual agreements, Service Level Agreements (SLAs), business relationships, or past perceptions of dCDN quality. As an alternative, such additional information could also be explicitly included with the given footprint.

It is reasonable to assume that a significant part of the actual footprint advertisement will occur out-of-band, prior to any CDNI FCI advertisement, with footprints defined in contractual agreements between participating CDNs. 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 will not be too detailed.

Given that a large part of the footprint advertisement will be defined 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 upon 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 previously agreed to serve in a contract with a uCDN.

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

- o A 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 A 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 footprint that is defined by coverage/reachability, or
- o capabilities of resources, such as delivery protocol, redirection mode, and metadata, which apply to a footprint that is defined by resources.

Resource footprint types are more specific than coverage/reachability footprint types, where the actual coverage and reachability are extrapolated from the resource location (e.g., a netmask applied to a resource IP address to derive an IP prefix). The specific methods

for extrapolating coverage/reachability from the resource location are beyond the scope of this document. In the degenerate case, the resource address could be specified as a coverage/reachability footprint type, in which case no extrapolation is necessary. Resource footprint types could expose the internal structure of a CDN; this could be undesirable. As such, the resource footprint types 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 footprint types narrows the dCDN's 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. However, making incremental updates 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 indicate 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 service 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 a footprint and capabilities are tied together and cannot be interpreted independently of 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 advertisement 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 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) are beyond the scope of 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, 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 that, in many cases, is absolutely necessary for deciding 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 is capable 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 also occur out-of-band, prior to any CDNI FCI advertisement, with capabilities defined in contractual agreements between participating CDNs. 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 to 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 is 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 reached regarding 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. FCI object extensibility is necessary to support future capabilities, as well as a generic protocol for conveying any capability information (e.g., with common encoding, error handling, and security mechanisms; further requirements for the CDNI FCI are listed in [RFC7337]).

Acknowledgments

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 Stiemerling 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 for the text.

Authors' Addresses

Jan Seedorf
HFT Stuttgart - University of Applied Sciences Stuttgart
Schellingstrasse 24
Stuttgart 70174
Germany

Phone: +49-0711-8926-2801
Email: jan.seedorf@hft-stuttgart.de

Jon Peterson
NeuStar
1800 Sutter St. Suite 570
Concord, CA 94520
United States of America

Email: jon.peterson@neustar.biz

Stefano Previdi
Cisco Systems
Via Del Serafico 200
Rome 0144
Italy

Email: sprevidi@cisco.com

Ray van Brandenburg
TNO
Anna van Buerenplein 1
The Hague 2595DA
The Netherlands

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

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