

CDNI WG
Internet-Draft
Intended status: Standards Track
Expires: March 23, 2013

S. Previdi, Ed.
F. Le Faucheur
J. Medved
Cisco Systems, Inc.
A. Guillou
SFR
September 19, 2012

CDNI Footprint Advertisement
draft-previdi-cdni-footprint-advertisement-02

Abstract

This document describes the use of BGP for Content Delivery Networks (CDNs) in order to advertise information about footprint and connectivity to footprint in the context of CDNI.

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 23, 2013.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents

(<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	4
2.	Terminology	4
3.	CDNI Mesh and MP-BGP	5
4.	CDNI Information	6
4.1.	Footprint	6
4.2.	Footprint Elements (FPE)	6
4.3.	Footprint Reachability (FPR)	7
4.4.	CDN Capability Information (CAP)	7
5.	CDNI Functional Components	8
5.1.	CDNI Databases	8
5.1.1.	CDNI Footprint Elements Database	8
5.1.2.	CDNI Footprint Reachability Database	9
5.1.3.	CDNI Capability Database	9
5.2.	CDNI MP-BGP Messages	9
5.2.1.	CDNI MP-BGP Footprint Element Advertisement	9
5.2.2.	CDNI MP-BGP Footprint Reachability Advertisement	10
5.2.3.	CDNI Capability Advertisement	10
5.3.	CDNI Information Elements (Attributes)	10
5.3.1.	CDN Identifier	10
5.3.2.	Footprint Element Identifier	10
5.3.3.	Reachable Footprint Element	11
5.3.4.	Origin_AS_PATH	11
5.3.5.	CDN Capabilities	11
5.4.	CDNI MP-BGP NLRIs	11
6.	CDNI Example	12
6.1.	Topology	12
6.2.	CDN2 Footprint Element Advertisements	12
6.3.	Footprint Element Database in each CDN	13
6.4.	Footprint Reachability Advertisements	14
6.5.	Capability Advertisements	16
6.6.	FPE, FPR and CAP Databases in uCDN	16
6.7.	Request Routing for user 2.2.2.2	17
6.8.	Request Routing for user 1.1.1.1	18
7.	CDNI MP-BGP Encodings	19
7.1.	CDNI Attributes	19
7.1.1.	CDNI Identifier	19
7.1.2.	FPE Identifier	19

7.1.3.	Origin_AS_PATH	19
7.1.4.	Capabilities	20
7.2.	CDNI Messages	20
7.2.1.	FPE Advertisement	20
7.2.2.	FPR Advertisement	21
7.2.3.	CAP Advertisement	21
7.3.	CDNI NLRI	22
7.3.1.	CDNI-NLRI Type 1: FPE Advertisement	22
7.3.2.	CDNI-NLRI Type 2: FPR Advertisement	23
7.3.3.	CDNI-NLRI Type 3: CAP Advertisement	23
7.3.4.	NLRI Encoding	24
8.	Compliance with CDNI Requirements	25
9.	IANA Considerations	25
10.	Security Considerations	25
11.	Acknowledgements	25
12.	References	25
12.1.	Normative References	25
12.2.	Informative References	26
	Authors' Addresses	26

1. Introduction

The IETF CDN Interconnection (CDNI) Working Group is chartered to develop specifications for interconnecting standalone CDNs so that their collective CDN footprint can be leveraged for the end-to-end delivery of content from Content Service Providers (CSPs) to End Users.

[I-D.jenkins-cdni-problem-statement] outlines the problem area that the CDNI working group is chartered to address.

[[I-D.bertrand-cdni-use-cases](#)] discusses the use cases for CDN Interconnection and [[I-D.davie-cdni-framework](#)] discusses the technology framework for the CDNI solution and interfaces.

When an upstream CDN (uCDN) receives a request from a user, it has to determine what is the downstream CDN (dCDN) to which the request is to be redirected. This CDN selection decision can take into account various criteria such as administrative preferences (for example based on the commercial arrangements between the uCDN and candidate dCDNs including associated request handling costs) and/or such as whether candidate CDNs have caches that are topologically close to the user and capable of handling that request. Therefore, as discussed in section "Dynamic Footprint Discovery" of [[I-D.davie-cdni-framework](#)], there are situations where being able to dynamically discover the set of requests that a given dCDN is willing and able to serve is beneficial. As also discussed in [[I-D.davie-cdni-framework](#)], this information could be potentially provided by the dCDN in response to a query by the uCDN, or the information (or its changes) could be spontaneously advertised by the dCDN.

The proposal outlined in this document makes use of Multiprotocol-BGP (MP-BGP [[RFC4760](#)]) in order for CDNs and/or ISPs to advertise their footprint information as well as for CDNs to advertise their connectivity to these footprints. In addition CDNs use MP-BGP advertisements to represent their interconnectivity.

2. Terminology

The following terminology is used in this document:

CDNI Mesh: the set of CDNs participating into CDNI and using MP-BGP sessions between them in accordance with the approach defined in this document. A CDNI Mesh has no requirements in terms of topology, i.e.: the mesh can be partial, full or hierarchical.

Footprint (FP): The exhaustive set of Prefixes a CDN is willing to serve.

Footprint Element (FPE): Arbitrary set of prefixes with attributes. They can be implicit (inferred from BGP) or explicit (advertised).

Implicit Footprint Element Advertisement (FPE Implicit Advertisement): Footprint Element information derived from the BGP database.

Footprint Element Advertisement (FPE-Adv): MP-BGP Message used by a CDN in order to advertise or withdrawn Footprint Elements and their attributes.

Footprint Reachability Advertisement (FPR-Adv): MP-BGP Message used by a CDN in order to advertise or withdrawn Footprint reachability information.

Capability Advertisement (CAP-Adv): MP-BGP Message used by a CDN in order to advertise or withdrawn capability information.

3. CDNI Mesh and MP-BGP

CDNI enables CDNs to communicate in order to deliver content in a collaborative mode.

CDNI Mesh will make use of Multiprotocol-BGP (MP-BGP [[RFC4760](#)]) for the exchange of footprint, reachability and capability information.

We define a new Sub Address Family (CDNI-SAFI, TBD) for Address Families (AF) 1 (IPv4) and 2 (IPv6) and a set of new Network Layer Reachability Information (NLRI) that will carry either FPE, FPR or CAP advertisements. The NLRI will have a NLRI type determining the content (i.e.: FPE, FPR or CAP) so to distinguish advertisements.

The advantage of using a separate address family is to isolate CDNI information from regular BGP-4 Internet information so to not compromise in any way the security and reliability of the current BGP information exchange used for IP network layer routing.

The advantage of having separate FPE, FPR and CAP information is that a CDN needs not originate (and update) FPE information each time there's a change in the way a CDN is connected to other CDNs. E.g.: if the complete Footprint information was to be exchanged between CDNs, it would consist of a very large amount of prefixes advertised (and re-advertised) each time a CDNI interconnection changes in the

CDNI Mesh.

When an existing CDNI connection is removed or when a new connection between two CDNs is established, the only advertisements that need to be updated are the ones concerning the FPR.

While footprint information is expected to be relatively stable, the CDNI Mesh (i.e.: the connectivity between CDNs) and the reachability from CDN to footprints may be impacted by network events. Also, the connectivity between CDNs may be affected by the CDN selection policy which may be modified relatively frequently.

CDNI FPR advertisements allow the CDNI Mesh to scale by adapting easily to topology changes. In fact, just a few number of FPR advertisements are used by each CDN which makes the CDNI MP-BGP scheme very scalable.

4. CDNI Information

CDNI Information includes:

Footprint (FP).

Footprint Element (FPE).

Footprint Reachability (FPR).

CDN Capabilities (CAP).

4.1. Footprint

Footprint represents the exhaustive set of prefixes a CDN is willing to serve. The set of prefixes may belong to one or more Autonomous Systems or may include subset of prefixes of one or more Autonomous Systems. In other words, there is no strict relationship between an Autonomous System and the CDN Footprint.

Example: if CDN-A delivers content to ISP-A users, then CDN-A footprint consists of all prefixes owned and connected to ISP-A.

4.2. Footprint Elements (FPE)

Footprint Elements refers to an arbitrary set of prefixes part of a CDN Footprint.

FPEs can be inferred from the BGP-4 Internet database. It is assumed that the CDN will have a BGP-4 feed with Internet prefixes that are

necessary in order for the CDNI Mesh to operate and from which it will be able to derive FPEs.

The CDN will maintain a FPE Database that is isolated from the regular BGP Internet database. The CDNI FPE database uses the CDNI-SAFI (TBD) MP-BGP address family.

FPEs can be implicit (inferred from BGP) or explicit (advertised).

Example 1: if CDN-A is willing to serve Autonomous System A and Autonomous System B users, then two footprint elements will be inferred from the BGP database in all CDNs: FPE-ASA and FPE-ASB. Then, CDN-A will advertise its reachability to FPEs FPE-ASA and FPE-ASB. In this example FPEs are inferred by all CDNs by just looking at the BGP-4 Internet Database.

Example 2: if CDN-B is willing to serve Autonomous System A users and a subset of Autonomous System B users, then CDN-B Footprint consist of two FPEs: Autonomous System A and the set of prefixes representing the subset of Autonomous System B users. Autonomous System 'A' FPE is inferred from BGP-4 Database while the subset of Autonomous System B users is an FPE that will be explicitly advertised by CDN-B to the CDNI Mesh.

The CDNI FPE Database includes FPEs inferred from the BGP-4 (Internet) database as well as FPE Advertisements originated by the CDN and/or received from other CDNI Mesh members.

4.3. Footprint Reachability (FPR)

Footprint Reachability refers to the way a CDN can reach to one or more FPEs. FPR consists of advertisements originated by each CDN for each of its FPE and advertised to the CDNI Mesh (through MP-BGP Messages).

FPR Information is stored in the CDNI FPR Database which contains the information originated by the CDN and the information received from the other CDNs in the CDNI Mesh.

4.4. CDN Capability Information (CAP)

Each CDN advertises to the CDNI Mesh the set of its capabilities. The CDN originates a CDN Capabilities message (CAP) message containing the attributes describing such capabilities.

Each CDN MUST originate a Capabilities message and each CDN will store in the CDNI Capabilities Database the set of Capabilities messages received from other members of the CDNI Mesh.

5. CDNI Functional Components

This section describes the CDNI functional components (databases, messages, information sets) stored and exchanged between CDNs and related to CDN Footprints and Capabilities. This document proposes the use of Multiprotocol-BGP for exchanging the CDNI information. The details of MP-BGP encodings are described in [Section 7](#).

5.1. CDNI Databases

This section describes the set of databases defined for the propagation of Footprint and Capability information across the CDNI Mesh. The goal of this section is to explain the different information elements that are required in order to properly determine FPE, FPR and CAP information. This section is not to be intended as an implementation description. An implementation may combine or merge all data into single or multiple databases.

CDNI Footprint Elements Database contains FPEs that are either derived from BGP-4 Internet table or received from other CDN through CDNI Footprint Elements Advertisements.

CDNI Footprint Reachability Database contains the advertisements originated by the CDN and received from other CDNs and that describe how each CDN can reach its Footprint Elements.

CDNI Capability Database contains the set of CDN capabilities advertised by each CDN.

5.1.1. CDNI Footprint Elements Database

FPE Database contains the FPEs that a CDN has inferred from the BGP-4 Internet database as well as FPEs that have been explicitly advertised by other CDNs.

FPE Database includes prefixes and their attributes as known in the BGP-4 Internet Database (e.g.: prefix, AS_PATH, MED, Communities, etc) in addition to CDNI specific attributes that are defined in following sections.

FPE Database also contains the FPEs that have been explicitly advertised by remote CDNs.

It is expected that for most of CDNs, the majority of the FPE information is inferred from BGP-4 Internet Database so to reduce significantly the advertisement of FPEs.

5.1.2. CDNI Footprint Reachability Database

FPR Database contains the set of FPEs each CDN claimed reachability to. This includes:

The set of FPEs this CDN can reach.

The set of FPEs other CDNs have advertised their reachability to.

The FPR Database is populated by the FPR Advertisement messages originated and sent by each CDN participating into the CDNI Mesh.

5.1.3. CDNI Capability Database

Capability (CAP) Database contains the set of capabilities advertised by each CDN.

The CAP Database is populated by the CAP Advertisement messages originated and sent by each CDN participating into the CDNI Mesh.

5.2. CDNI MP-BGP Messages

We define three new MP-BGP messages for the advertisement of CDNI Address Family NLRIs:

Footprint Element Advertisement (FPE Advertisement) describing an FPEs and its attributes.

Footprint Reachability Advertisement (FPR Advertisement) describing how a CDN can reach one or more FPEs.

Capability Advertisement (CAP Advertisement) describing CDN capabilities.

The propagation of CDNI Messages throughout the CDNI Mesh is done according to standard BGP rules and the inter-CDN connectivity will be reflected in the BGP attributes, e.g.: AS_PATH will describe the different CDNs the advertisement traversed during its propagation in the CDNI Mesh thus describing the inter-CDN connectivity.

5.2.1. CDNI MP-BGP Footprint Element Advertisement

FPE Advertisement is originated by a CDN willing to group prefixes it can reach under a unique identifier (identifying the FPE). The message is originated and sent out by the CDN and propagated throughout the CDNI Mesh according to MP-BGP propagation rules.

The FPE Advertisement includes CDNI NLRI, representing the FPE

address space, and attributes among which the identifier of the FPE.

5.2.2. CDNI MP-BGP Footprint Reachability Advertisement

In the CDNI Mesh, each CDN MUST advertise the different Footprint Elements it reaches. The CDN maintains a FPE Reachability Database with entries describing its connectivity to Footprints Elements.

The propagation of CDN Connectivity advertisements throughout the CDNI Mesh is done according to standard BGP rules and the inter-CDN connectivity will be reflected in the BGP attributes (e.g.: AS_PATH will describe the different CDNs the advertisement traversed during its propagation in the CDNI Mesh thus describing the inter-CDN connectivity).

5.2.3. CDNI Capability Advertisement

Each CDN MUST advertise information describing its capabilities. This is done in the CAP dvertisement message.

5.3. CDNI Information Elements (Attributes)

This section describes the different information elements contained in the CDNI Databases.

5.3.1. CDN Identifier

CDN_Identifier uniquely identifies the CDN within the CDNI Mesh. CDN Identifier is exchanged inside the CDNI (MP-BGP) messages between CDNs.

5.3.2. Footprint Element Identifier

The role of the Footprint Element Identifier is to group all prefixes part of the same FPE under a unique identifier. This allows a CDN to claim reachability to the FPE by just specifying the FPE Identifier rather than each individual prefix of the footprint.

When inferring the CDNI footprint information from the regular BGP-4 Internet database, the FPE Identifier is derived from the first ASN in the AS_PATH of the prefix, i.e.: the AS of origin of the prefix.

FPE Identifiers are present in both FPE Advertisement and FPR Advertisement. In the FPE Advertisement it is used to arbitrarily group prefixes under the same Identifier while in the FPR Advertisement, the FPE Identifier is used in order to describe which FPEs a CDN can reach.

5.3.3. Reachable Footprint Element

Reachable FPE contains the FPE Identifier of each FPE the CDN claim reachability to. Multiple instances of Reachable FPE are allowed in the same FPR Advertisement message. FPE Identifier is assigned to the FPE by the CDN advertising the FPE (in the FPE Advertisement Message). FPE Identifier contains an AS number when the FPE has been inferred from the BGP database.

5.3.4. Origin_AS_PATH

This attribute contains the prefix AS_PATH value that is present in the BGP-4 Internet Database of the CDN originating the FPR advertisement.

Origin_AS_PATH aims to describe how the CDN can reach a given FPE and is present only in FPR advertisements.

5.3.5. CDN Capabilities

Capabilities are expressed in a set of attributes the CDN inserts in the Capability Advertisement message.

5.4. CDNI MP-BGP NLRI

We define three types of CDNI NLRI: FPE-NLRI, FPR-NLRI and CAP-NLRI

When referred to FPE, the NLRI contains either an IPv4 or IPv6 prefix.

In some cases, a given FPE prefix is reachable by multiple CDNs. MP-BGP, according to BGP selection rule, will allow to select only one FPE and therefore information about FPE being multi-reachability through different CDNs may be lost. In order to prevent this, the CDNI NLRI contains a CDNI FPE Distinguisher so to make two identical prefixes different from a BGP selection perspective. The behavior is similar to what defined in [[RFC4364](#)].

FPR NLRI consists of one or more FPE-Identifiers representing the FPEs the CDN can reach.


CAP-NLRI contains a IPv4 or IPv6 prefix (called the CAP-Prefix) belonging to the CDN address space originating the CAP Advertisement Message. Capabilities will be advertised using attributes to the CAP-prefix.

6. CDNI Example

This section illustrates an example of CDNI Mesh with the set of databases and messages that are exchanged between CDNs.

6.1. Topology

The following topology is used:



CDN2, CDN3 and uCDN participate into the CDNI Mesh.

CDNI Mesh includes following MP-BGP Sessions: CDN2-CDN3 and CDN3-uCDN. Note that there's no MP-BGP session between CDN2 and uCDN.

AS100 and AS400 have no CDN.

AS100 advertises all prefixes to AS300 and AS400.

AS100 advertises only prefix 1.1.1.0/24 and 3.3.3.0/24 to AS200.

6.2. CDN2 Footprint Element Advertisements

CDN2 originates following Footprint Elements Advertisements:


```
Prefix CDN2:1.1.1.0/24
  AS_PATH: {CDN2}
  Footprint-Element Identifier: FP-CDN2
```

```
Prefix CDN2:3.3.3.0/24
  AS_PATH: {CDN2}
  Footprint-Element Identifier: FP-CDN2
```

where:

Prefix CDN2:1.1.1.0/24 and Prefix CDN2:3.3.3.0/24 are the two prefixes advertised with the CDNI FPE Distinguisher corresponding to CDN2. The CDNI FPE Distinguisher is required so to distinguish between possible identical prefixes advertised by other CDNs.

AS_PATH contains the set of CDNs the update has traversed. It is mainly used for loop prevention in BGP in the CDNI Mesh.

Note: the AS numbers in the AS_PATH represent the CDNs, not the BGP autonomous systems.

Footprint-Element Identifier contains the identifier of the Footprint-Element and it is used by CDN2 when originating the Footprint Reachability Advertisement.

Note that uCDN and CDN3 do NOT originate any Footprint Element Message.

6.3. Footprint Element Database in each CDN

Footprint Elements Database is constructed by inferring all info from BGP database plus the Footprint Elements Advertisements received from within the CDNI Mesh.

In CDN2 the Footprint Elements Database is as follows:

```
1.1.1.0/24; AS_PATH: {100}      ===\
2.2.2.0/24; AS_PATH: {300, 100}  ====> inferred from
3.3.3.0/24; AS_PATH: {100}      ===/   BGP table
```

```
CDN2:1.1.1.0/24
  AS_PATH: {}
  Footprint-Element Identifier: FP-CDN2
```

```
CDN2:3.3.3.0/24
  AS_PATH: {}
  Footprint-Element Identifier: FP-CDN2
```


In CDN3 the Footprint Elements Database is as follows:

```
1.1.1.0/24; AS_PATH: {100}      ===\  
2.2.2.0/24; AS_PATH: {100}      =====> inferred from  
3.3.3.0/24; AS_PATH: {100}      ===/      BGP table
```

```
CDN2:1.1.1.0/24  
  AS_PATH: {CDN2}  
  Footprint-Elements Identifier: FP-CDN2
```

```
CDN2:3.3.3.0/24  
  AS_PATH: {CDN2}  
  Footprint-Element Identifier: FP-CDN2
```

In uCDN the Footprint Elements Database is as follows:

```
1.1.1.0/24; AS_PATH: {<internet>, 400, 100} ===\  inferred  
2.2.2.0/24; AS_PATH: {<internet>, 400, 100} =====>  from  
3.3.3.0/24; AS_PATH: {<internet>, 400, 100} ===/  BGP table
```

```
CDN2:1.1.1.0/24  
  AS_PATH: {CDN3, CDN2}  
  Footprint-Element Identifier: FP-CDN2
```

```
CDN2:3.3.3.0/24  
  AS_PATH: {CDN3, CDN2}  
  Footprint-Element Identifier: FP-CDN2
```

The Footprint-Elements Database contains prefixes inferred from the BGP database and the prefixes received through Footprint-Element Advertisements.

Footprint-Element prefixes derived from BGP database do not require a CDNI FPE Distinguisher. The way an implementation stores inferred FPE information (with or without CDNI FPE Distinguisher is out of the scope of this document.

The prefixes received by FPE Advertisements need a CDNI FPE Distinguisher because multiple CDNs may advertise the same set of prefixes and we need to preserve all from selection rules.

6.4. Footprint Reachability Advertisements

Each CDN will advertise its reachability to Footprint Elements:

CDN2 Footprint Reachability Advertisement:

CDN-Identifier: CDN2

AS_PATH: <will contain the AS_PATH representing the path
taken by Footprint-Reachability advertisement in
the CDNI Mesh>

Reachable Footprint-Element: AS200

Origin_AS_PATH: {}

Reachable Footprint-Element: FP-CDN2

Origin_AS_PATH: {100}

Where:

Origin_AS_PATH describes the BGP AS_PATHs as known in CDN2.

CDN2 advertises reachability to two sets of prefixes: one consists of all prefixes in AS200 and another set consists of prefixes he received from AS100 (and that he previously sent with Footprint-Element Advertisements).

CDN3 Footprint Reachability Advertisement:

CDN-Identifier: CDN3

AS_PATH: <will contain the AS_PATH representing the path
taken by Footprint-Reachability advertisement in
the CDNI Mesh>

Reachable Footprint-Element: AS300

Origin_AS_PATH: {}

Reachable Footprint-Element: AS100

Origin_AS_PATH: {200, 100}

Where:

Origin_AS_PATH describes the BGP AS_PATHs as known in CDN3.

CDN3 advertises reachability to two sets of prefixes: AS100 and AS300.

Each CDN advertises:

CDNI Identifier.

Set of reachable Footprint-Elements (i.e.: a set of FPE Identifiers) with their corresponding Origin_AS_PATH.

The AS_PATH representing the BGP connectivity within the CDNI Mesh.

6.5. Capability Advertisements

Each CDN originates a CAP Advertisement messages carrying information about its capabilities. Each CDN will advertise:

A prefix owned by the CDN (i.e.: part of the CDN address space).

Its CDN Identifier.

A set of attributes describing the CDN capabilities (encoded as BGP Extended Community Attributes).

CDN2 Capability Advertisement:

CDN-Identifier: CDN2

Communities (std/ext): representing CDN2 capabilities

AS_PATH: <will contain the AS_PATH representing the path taken by Footprint-Reachability advertisement in the CDNI Mesh>

NLRI: CDN2-Prefix

CDN3 Capability Advertisement:

CDN-Identifier: CDN3

Communities (std/ext): representing CDN capabilities

AS_PATH: <will contain the AS_PATH representing the path taken by Footprint-Reachability advertisement in the CDNI Mesh>

NLRI: CDN3-Prefix

6.6. FPE, FPR and CAP Databases in uCDN

uCDN has the two following databases:

Footprint-Elements Database:

1.1.1.0/24

AS_PATH: {<internet>, 400, 100} ==\

2.2.2.0/24 ===\

AS_PATH: {<internet>, 400, 100} =====> inferred from

3.3.3.0/24 ===/ BGP table

AS_PATH: {<internet>, 400, 100} ==/

CDN2:1.1.1.0/24

AS_PATH: {CDN3, CDN2}

Footprint-Element Identifier: FP-CDN2

CDN2:3.3.3.0/24

AS_PATH: {CDN3, CDN2}

Footprint-Element Identifier: FP-CDN2

Footprint-Reachability Database:

CDN-Identifier: CDN2

AS-PATH: {300, 200}

Reachable Footprint-Element: AS200

Origin_AS_PATH: {}

Reachable Footprint-Element: FP-CDN2

Origin_AS_PATH: {100}

CDN-Identifier: CDN3

AS-PATH: {300}

Reachable Footprint-Element: AS300

Origin_AS_PATH: {}

Reachable Footprint-Element: AS100

Origin_AS_PATH: {200, 100}

Capability Database:

CDN-Identifier: CDN2

Communities (std/ext): representing CDN2 capabilities

AS_PATH: <will contain the AS_PATH representing the path
taken by Footprint-Reachability advertisement in
the CDNI Mesh>

NLRI: CDN2-Prefix

CDN-Identifier: CDN3

Communities (std/ext): representing CDN capabilities

AS_PATH: <will contain the AS_PATH representing the path
taken by Footprint-Reachability advertisement in
the CDNI Mesh>

NLRI: CDN3-Prefix

The Footprint-Reachability and Capabilities databases also contains the AS_PATH representing the MP-BGP connectivity in the CDNI Mesh and the path taken by Footprint-Reachability Advertisements. It is anyway mandatory in BGP for message loop prevention.

6.7. Request Routing for user 2.2.2.2

The following workflow is used for a request coming from user 2.2.2.2:

User with address 2.2.2.2 sends a content request to uCDN.

uCDN does a lookup in Footprint-Elements Database for address 2.2.2.2 and finds following entry:

2.2.2.0/24

AS_PATH: {<internet>, 400, 100}

uCDN does a lookup in Footprint-Reachability Database and look for a CDN claiming connectivity to AS100. It finds following entry:

CDN-Identifier: CDN3

AS_PATH: {300}

Reachable Footprint-Element: AS300

Origin_AS_PATH: {}

Reachable Footprint-Element: AS100

Origin_AS_PATH: {200, 100}

uCDN can select CDN3 as downstream CDN.

6.8. Request Routing for user 1.1.1.1

The following workflow is used for a request coming from user 1.1.1.1:

User with address 1.1.1.1 sends a content request to uCDN.

uCDN does a lookup in Footprint-Elements Database for address 1.1.1.1 and finds following entries:

1.1.1.0/24

AS_PATH: {<internet>, 400, 100}

and

CDN2:1.1.1.0/24

AS_PATH: {CDN3, CDN2}

Footprint-Element Identifier: FP-CDN2

uCDN prefers the entry originated by a Footprint-Element Advertisement and looks into the Footprint-Reachability Database for a CDN claiming connectivity to Footprint-Element Identifier: FP-CDN2. It finds following entries:

CDN-Identifier: CDN2

AS_PATH: {300, 200}

Reachable Footprint-Element: AS200

Origin_AS_PATH: {}

Reachable Footprint-Element: FP-CDN2

Origin_AS_PATH: {100}

uCDN can select CDN2 as downstream CDN

7. CDNI MP-BGP Encodings

This section describes the CDNI MP-BGP messages with their attributes and NLRI encodings details.

7.1. CDNI Attributes

CDNI uses standard BGP attributes in both FPE and FPR advertisements in addition to the following newly defined attributes:

FPE_Identifier is present in FPE advertisements (when identifying an FPE as a set of prefixes) and FPR advertisements (when identifying a reachable FPE).

CDN_Identifier is present in both FPE and FPR advertisements).

Origin_AS_PATH is present in FPR advertisements.

CDN Capabilities is present in CAP advertisements.

7.1.1. CDNI Identifier

CDN_Identifier is a mandatory transitive attribute consisting of an extended community (type TBD) uniquely identifying the CDN within the CDNI Mesh. CDNI_Identifier must be unique within the CDNI Mesh. The recommended method is to use (as for the FPE_Identifier) the AS or IP Address format as described in[RFC4360].

7.1.2. FPE Identifier

FPE Identifier is a mandatory and transitive attribute consisting of an extended community (type TBD) containing a Footprint Element identifier. The FPE_Identifier uniquely identifies the FPE within the CDNI Mesh. The format of the FPE_Identifier encodings should ensure uniqueness of values. The recommended method for encoding FPE_Identifier is to use the format as defined in [RFC4360](#) and use either the AS or an IP address owned by the CDN. This method, and the extended community format, allows a CDN to create multiple unique FPE_Identifiers.

7.1.3. Origin_AS_PATH

Origin_AS_PATH is an optional transitive attribute having the same format as AS_PATH attribute and containing the AS_PATH value of the Footprint Element known in the CDN originating the FPR advertisement.

Origin_AS_PATH must be left unchanged and has the same format as AS_PATH BGP attribute as defined in [[RFC4271](#)].

7.1.4. Capabilities

Capabilities is a mandatory and transitive attribute consisting of an extended community (type TBD) containing a value representing a given capability (TBD). More than one Capability attribute is allowed so to allow a CDN to advertise multiple capabilities in the same message. The encodings of Extended Communities is defined in: [[RFC4360](#)].

7.2. CDNI Messages

Three messages are used: FPE Advertisement, FPR Advertisement and CAP Advertisement.

7.2.1. FPE Advertisement

Footprint Element Advertisement message is used by a CDN in order to advertise a Footprint Element that can't be inferred from the BGP-4 Internet Database and that needs explicit advertisement.

The FPE Advertisement consists of one or more prefixes that a CDN reaches. The FPE Advertisement scope is not to describe how the CDN can reach the FPE but rather what are the prefixes forming the FPE. The FPE Advertisement message should contain following information:

CDN_Identifier: uniquely identifies the CDN originating the message. It is mandatory and transitive attribute.
CDN_Identifier format is described in [Section 7.1](#).

Footprint-Element Identifier: a number uniquely identifying the FPE as defined in [Section 7.1](#). FPE-Identifier is encoded as an extended community. It is mandatory and transitive.

Standard BGP attributes such as AS_PATH, Communities, Local_Preference, MED, etc. are also used in the FPE Advertisement Messages.

NLRI: the set of prefixes (address and mask) of the FPE.

A CDNI Distinguisher is prepended to each prefix of the FPE. The CDNI Distinguisher allows multiple CDNs to advertise the same prefix in their FPE Advertisement. Each CDN must use a unique number. The recommended method for assigning RDs is to use the CDN AS number or an IP address owned by the CDN.

7.2.2. FPR Advertisement

Footprint Reachability Advertisement message is used by a CDN willing to advertise the FPEs it can reach. The FPR Advertisement consists of a set of FPE-Identifiers and their attributes. Following information should be present in the FPR Advertisement message:

CDN_Identifier: uniquely identifies the CDN originating the message. It is mandatory and transitive attribute.

CDN_Identifier format is described in [Section 7.1](#).

Reachable Footprint-Element: one or more AS numbers and/or FPE_Identifier representing the reachable FPEs. Reachable_FPE is encoded as a set of extended communities. At least one must be present in the message and the attribute is transitive. FPEs that have been inferred from the BGP4 Internet database are identified through their AS number and FPEs that have been explicitly advertised are identified by their FPE_Identifier.

Standard BGP attributes such as AS_PATH, Communities, Local_Preference, MED, etc. are also used in the FPR Advertisement Messages.

Origin_AS_PATH. This attribute is related to an FPE and contains the value of the BGP AS_PATH of the FPE in the CDN originating the FPR message. When used, it implies all prefixes of the same FPE share the same BGP AS_PATH in the CDN originating the FPR message.

NLRI: Reachable Footprint-Element: one or more FPE Identifiers representing the reachable FPEs. Reachable FPE is encoded in the CDNI NLRI. At least one must be present in the message. FPEs that have been inferred from the BGP4 Internet database are identified through their AS number and FPEs that have been explicitly advertised are identified by their FPE Identifier.

7.2.3. CAP Advertisement

Capability Advertisement message is used by a CDN in order to advertise its capabilities. Capabilities are encoded as a set of BGP Extended Communities. The CAP Advertisement Message contains:

A CDN_Identifier: uniquely identifies the CDN originating the message. It is mandatory and transitive attribute.

CDN_Identifier format is described in [Section 7.1](#).

One or more Extended Communities attributes describing capabilities.

Standard BGP attributes such as AS_PATH, Communities, Local_Preference, MED, etc. are also used in the CAP Advertisement Messages for policy reasons.

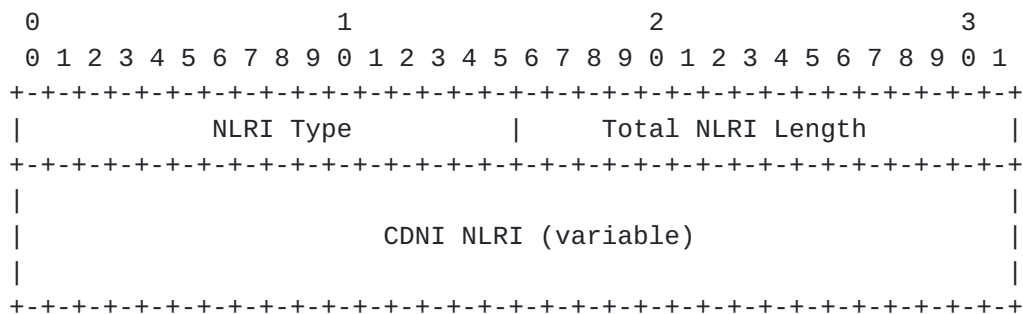
A CDN Prefix: a NLRI encoded prefix (IPv4 or IPv6) owned by the CDN (i.e.: a prefix that is part of the address space of the CDN).

7.3. CDNI NLRI

Multiprotocol-BGP defines two attributes, MP_REACH and MP_UNREACH, that are BGP's containers for carrying opaque information. Each CDNI NLRI describes either a FPE or a FPR.

All FPE and FPR information SHALL be encoded using AFI 1 or 2 (ipv4 or ipv6) with a TBD (CDNI) SAFI.

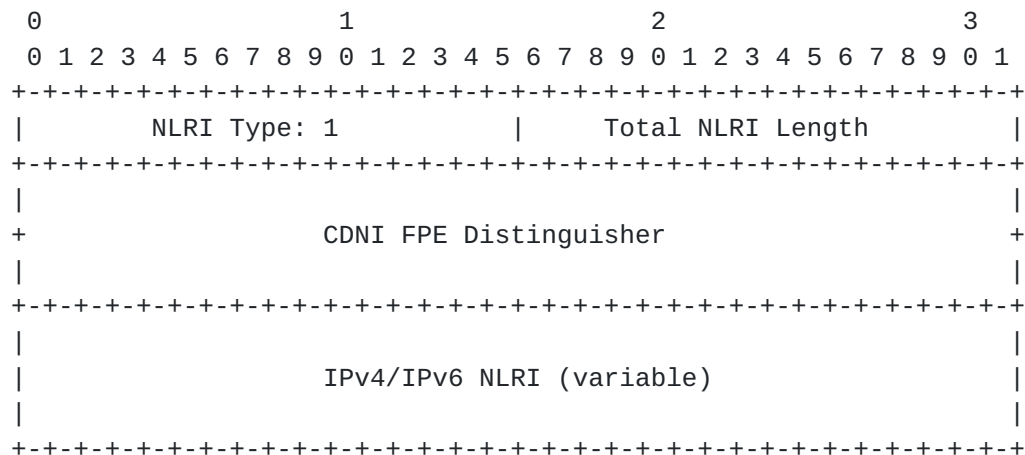
In order for two BGP speakers to exchange CDNI NLRI, they MUST use BGP Capabilities Advertisement to ensure that they both are capable of properly processing CDNI NLRI. This is done as specified in [\[RFC4760\]](#), by using capability code 1 (multi-protocol BGP), with the AFI: 1 or 2, SAFI: TBD.



where:

NLRI Type: 1 (FPE)
 2 (FPR)
 3 (CAP)

7.3.1. CDNI-NLRI Type 1: FPE Advertisement



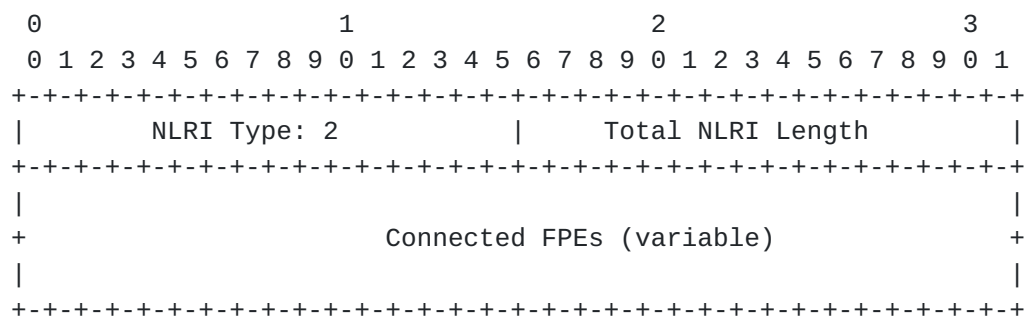
where:

NLRI Type: 1 specifies the FPE Advertisement

CDNI FPE Distinguisher: 8 octets

NLRI: IPv4 or IPv6 prefix

[7.3.2.](#) CDNI-NLRI Type 2: FPR Advertisement



where:

NLRI Type 2 specifies the FPR Advertisement

Connected FPEs: set of FPE Identifiers representing the FPEs the
CDN can reach.

[7.3.3.](#) CDNI-NLRI Type 3: CAP Advertisement

It has to be noted that a CDN may use multiple FPR messages for advertising reachability to multiple FPEs. In such cases multiple NLRI prefixes would be needed in order not to create collisions in the BGP selection process (that would select only one NLRI message among the ones having identical NLRI prefixes).

8. Compliance with CDNI Requirements

[I-D.ietf-cdni-requirements] outlines the requirements for the solution and interfaces to be specified by the CDNI working group. This section identifies the relevant requirements from that document and discusses compliance by the solution proposed in this document.

[Editor's Note: Text is to be added when requirements-03 is available. This needs to discuss the requirements labeled R27, R28, R29 and R30 as of requirements-02].

9. IANA Considerations

none.

10. Security Considerations

This draft does not affect the BGP security model.

11. Acknowledgements

The authors would like to thank Steven Luong, Manish Bhardwaj and Bruce Davie for their contributions.

12. References

12.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), January 2006.
- [RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", [RFC 4360](#), February 2006.
- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", [RFC 4364](#), February 2006.
- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", [RFC 4760](#), January 2007.

12.2. Informative References

- [I-D.bertrand-cdni-use-cases]
Bertrand, G., Stephan, E., Watson, G., Burbridge, T.,
Eardley, P., and K. Ma, "Use Cases for Content Delivery
Network Interconnection", [draft-bertrand-cdni-use-cases-02](#)
(work in progress), July 2011.
- [I-D.davie-cdni-framework]
Davie, B. and L. Peterson, "Framework for CDN
Interconnection", [draft-davie-cdni-framework-01](#) (work in
progress), October 2011.
- [I-D.ietf-cdni-requirements]
Leung, K. and Y. Lee, "Content Distribution Network
Interconnection (CDNI) Requirements",
[draft-ietf-cdni-requirements-03](#) (work in progress),
June 2012.
- [I-D.jenkins-cdni-problem-statement]
Niven-Jenkins, B., Faucheur, F., and N. Bitar, "Content
Distribution Network Interconnection (CDNI) Problem
Statement", [draft-jenkins-cdni-problem-statement-02](#) (work
in progress), March 2011.

Authors' Addresses

Stefano Previdi (editor)
Cisco Systems, Inc.
Via Del Serafico 200
Rome 00191
IT

Email: sprevidi@cisco.com

Francois Le Faucheur
Cisco Systems, Inc.
Greenside, 400 Avenue de Roumanille
Sophia Antipolis 06410
FR

Email: flefauch@cisco.com

Jan Medved
Cisco Systems, Inc.
3700 Cisco Way
SAN JOSE, CA 95134
US

Email: jmedved@cisco.com

Allan Guillou
SFR
40-42 Quai du Point du Jour
Boulogne-Billancourt 92659
FR

Email: allan.guillou@sfr.com

