

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: March 7, 2013

G. Bertrand
France Telecom - Orange
September 3, 2012

CDN Footprint Discovery
draft-bertrand-cdni-footprint-discovery-01

Abstract

Interconnected CDNs need to exchange information on the set of end-users to which they can deliver content. This information is commonly referred to as "CDN Footprint". This memo presents use cases for CDN Footprint Discovery in CDNI. It provides a survey of existing work on the subject and a set of additional requirements for controlling the exchange of Footprint information.

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 7, 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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1.	Terminology	3
2.	Introduction	4
2.1.	Abbreviations	4
3.	Use-Cases for Footprint Discovery	5
3.1.	Protocol Not Required	5
3.2.	Protocol Potentially Required	6
4.	Additional Requirements	6
5.	Survey of CDN Footprint Discovery	7
5.1.	Legacy Protocols	7
5.1.1.	Legacy BGP	8
5.1.2.	Legacy BGP Community Tag	8
5.2.	New Proposals	8
5.2.1.	BGP Extension for CDNI	8
5.2.2.	ALTO Footprint	8
6.	Survey of CDN Delivery Proximity Discovery	9
6.1.	BGP-TE	9
6.2.	BGP AIGP	9
6.3.	BGP Extension for CDNI	9
6.4.	ALTO	9
6.5.	Generic Capability Advertisement	10
7.	Synthesis	10
8.	IANA Considerations	11
9.	Security Considerations	11
10.	Acknowledgments	11
11.	References	11
11.1.	Normative References	11
11.2.	Informative References	12
	Author's Address	13

1. Terminology

We adopt the terminology described in [\[I-D.ietf-cdni-problem-statement\]](#) and [\[I-D.davie-cdni-framework\]](#), and extend it with the additional terms defined below.

Aggregate CDN Footprint: a set of User-Agent reachability information for which a CDN claims that it can deliver content in good conditions, by itself or through one of its dCDNs. The CDN Footprint aggregates information on the footprint of the dCDNs with whom a uCDN is interconnected.

High-Level CDN Footprint: the part of the footprint information that reflects rather static and business-level information. As an example, the failure of a Surrogate does not change the High-level CDN Footprint information but may change detailed information of an element of the CDN Footprint.

On-Net Footprint: a set of User-Agent reachability information for which a CDN claims that it can deliver content directly. For instance, a given Access CDN may assert that its On-Net CDN Footprint encompasses all end-users in AS 64496 and AS 64497.

CDN Delivery Proximity: Information on the network distance between a set of end-users in the CDN Footprint and the closest Surrogates of the considered CDN or of one of its dCDNs. Various metrics can be considered for defining this distance; examples of such metrics include the AS hop count or an accumulated Interior Gateway Protocol (IGP) metric.

CDN Footprint Discovery: discovery of information on CDN Footprint and CDN Delivery Proximity. CDN Footprint Discovery provides the information that enables a uCDN's Request Routing Interface to select the most appropriate dCDN for a given content request. CDN Footprint Discovery encompasses two different parts:

1. High-Level Footprint Discovery permits discovering groups of end-users/Surrogates and interconnection costs between them.
2. Detailed Footprint Discovery permits exchanging information that is subject to more scalability and confidentiality constraints. The level of information sharing must be tightly controlled.

CDN Footprint Discovery Interface: An interface that enables CDN Footprint Discovery. [Section 3](#) details the use cases for a CDN Footprint Discovery Interface.

2. Introduction

This memo presents use cases for CDN Footprint and CDN Delivery Proximity discovery in CDNI. It provides a survey of existing work on the subject and a set of additional requirements for controlling the exchange of Footprint information.

The reader should be familiar with the work of the CDNI WG:

- o CDNI problem statement [[I-D.ietf-cdni-problem-statement](#)] defines the problem area that the CDNI working group is chartered to address.
- o [[I-D.ietf-cdni-use-cases](#)] outlines real world use-cases for interconnecting CDNs. These use cases (e.g., "QoE and QoS Improvement") require the discovery of CDN Footprint information.
- o [[I-D.ietf-cdni-requirements](#)] specifies a set of requirements for CDN Footprint Discovery.

The present document describes:

- o The use cases for CDN Footprint Discovery ([Section 3](#)),
- o The requirements for CDN Footprint Discovery ([Section 4](#)),
- o A survey of Footprint Discovery protocols ([Section 5](#)).

2.1. Abbreviations

- o CDN: Content Delivery Network
- o CDNP: Content Delivery Network Provider
- o CSP: Content Service Provider
- o dCDN: downstream CDN
- o IGP: Interior Gateway Protocol
- o NSP: Network Service Provider
- o uCDN: upstream CDN

3. Use-Cases for Footprint Discovery

The present memo considers the use cases for a Footprint Discovery Protocol in a multi-CDN case [[I-D.ietf-cdni-use-cases](#)]. It does not consider mono-CDN issues.

[I-D.davie-cdni-framework] ([Section 3.5.](#)) describes "Dynamic Footprint Discovery" as a situation where "being able to dynamically discover the set of requests that a given dCDN is willing and able to serve is beneficial. For example, a CDN might at one time be able to serve a certain set of client IP prefixes, but that set might change over time due to changes in the topology and routing policies of the IP network." [[I-D.davie-cdni-framework](#)] provides an example where footprint information exchanges occur through the request routing interface and are triggered by an end-user's request (DNS resolution or content request).

In the present memo, we seek to clarify in which cases such Footprint Discovery through a protocol is required.

3.1. Protocol Not Required

In most cases, High-Level CDN Footprint Discovery does not require a protocol, as in the following examples cases:

- o High-Level CDN Footprint is Germany
- o High-Level CDN Footprint is AS 64496

However, the two following special cases deserve particular attention:

1. When the dCDNs' Footprints overlap,
2. When end-users outside the dCDNs' Footprints can request content.

In these two cases, the uCDN needs additional criteria than the dCDN's Footprint to select the "best" CDN. For instance, a static rule can be configured to select one of the available dCDNs. Alternatively, the uCDN may use dCDN's Delivery Proximity information to elect a delivering CDN.

The remainder of this memo focuses on cases, where either CDN Footprint is defined dynamically or uCDN requires dynamic Delivery Proximity information on dCDNs.

3.2. Protocol Potentially Required

In some cases, the uCDN needs dCDNs' Delivery Proximity information to determine which dCDN is the "best" to serve a given set of end-users. For instance, the "best" CDN can be defined as the one that has deployed Surrogates topologically the closest to the end-user (e.g., an Access CDN). This topological information is tightly related to the underlying networks: one may consider that a Surrogate is topologically far from the end-user if the path between these two entities crosses high cost links or many links. While information about the Surrogates location is rather confidential, abstract information about the path's cost between a set of end-users and the closest Surrogates in a CDN may be generic enough to avoid disclosing confidential information about the network.

CDN federations might involve several levels of CDN interconnection. In such scenarios, the CDN Footprint of a given CDN represents the aggregation of the CDN Footprint of all its own dCDNs. Therefore, some variations of the dCDNs' Footprint can result in variations of the CDN's Footprint. This example shows that the more levels of delegation we consider, the more dynamic the CDN Footprint information becomes. Nevertheless, in the medium term, complex deployments scenarios involving more than a few levels of delegations are unlikely, because of performance issues. Consequently, we consider that the High-Level CDN Footprint is rather static and remains valid for at least 24 hours.

Depending on the considered metric, the CDN Delivery Proximity may change rarely (e.g., AS hop count) or more frequently (e.g., accumulated IGP metric).

4. Additional Requirements

[[I-D.ietf-cdni-requirements](#)], already specifies two requirements related to Footprint Discovery: REQ-2 and REQ-3. We remind these two requirements below.

"REQ-2 [MED] The CDNI Request-Routing interface should allow the Downstream CDN to communicate to the Upstream CDN aggregate information to facilitate CDN selection during request routing, such as Downstream CDN capabilities, resources and affinities (i.e. Preferences or cost). This information could, for example, include:

- o supported content types and delivery protocols
- o footprint (e.g., layer-3 coverage)

- o a set of metrics/attributes (e.g., Streaming bandwidth, storage resources, distribution and delivery priority)
- o a set of affinities (e.g., Preferences, indication of distribution/delivery fees)
- o information to facilitate request redirection (e.g., Reachability information of Downstream CDN Request Routing system)."

"REQ-3 [MED] In the case of cascaded redirection, the CDNI Request-Routing interface shall allow the Downstream CDN to also include in the information communicated to the Upstream CDN, information on the capabilities, resources and affinities of CDNs to which the Downstream CDN may (in turn) redirect requests received by the Upstream CDN. In that case, the CDNI Request-Routing interface shall prevent looping of such information exchange."

We define additional requirements, specific to Footprint Discovery.

FPT-1 [MED] A uCDN must be able to discover CDN Footprint and CDN Delivery Proximity information about dCDNs.

FPT-2 [HIGH] A dCDN MUST be able to control what other CDNs can discover about its CDN Footprint and CDN Delivery Proximity.

We also clarify REQ-3:

FPT-3 [MED] A uCDN should not forward to any other CDN the Footprint and Delivery Proximity information that it has discovered about a dCDN without the explicit agreement of this dCDN.

Finally, the deployment of a Footprint Discovery Interface imposes some operational requirements. For instance, a Footprint Discovery protocol must not affect network stability and scalability. It must also be simple enough to avoid increasing the networks' operation complexity.

FPT-4 [HIGH] A Footprint Discovery protocol should not affect network stability and scalability.

5. Survey of CDN Footprint Discovery

5.1. Legacy Protocols

5.1.1. Legacy BGP

Consider a dCDN that claims its CDN Footprint covers AS 64496. BGP advertises AS-Path information that easily enables a uCDN to map end-users to the dCDN, basing on the end-users' IP address and on a mapping of IP addresses to AS numbers.

BGP also provides CDN Delivery Proximity information: for instance, a CDN listening to BGP advertisements is able to determine the AS-path length to the advertised prefixes. Note that BGP information might be collected in the network and provided to the CDN through another protocol rather than directly through BGP.

5.1.2. Legacy BGP Community Tag

The NSP may use part of the community tags carried by its legacy internal BGP to filter and gather the prefixes in stable groups (see section 5.1.7 of [[I-D.ietf-alto-deployments](#)]) that are then used by its internal CDN [[I-D.jenkins-alto-cdn-use-cases](#)] for fine-grained request routing based on these groups. A protocol (e.g., HTTP-based) can be used to advertise aggregated information on such groups rather than detailed information per prefix. This provides a simple way to aggregate information for scalability and confidentiality purposes.

An operator may consider that the grouping of prefixes into zones (the list of prefixes with a given community value) is confidential, as this grouping discloses information on the network's organization.

5.2. New Proposals

5.2.1. BGP Extension for CDNI

[I-D.previdi-cdni-footprint-advertisement] proposes a BGP-based mechanism to advertise connectivity information in the context of CDNI. It is based on the introduction of a CDN sub address Family (SAFI) and leverages BGP (extended) community tags. It uses Multiprotocol-BGP (MP-BGP [[RFC4760](#)]) in order for CDNs and/or ISPs to advertise their connectivity to footprints. A new NLRI is defined to carry CDN connectivity advertisements. In summary, the draft defines a "CDN-level BGP" to complement the legacy network-level BGP described in [Section 5.1.1](#) and [Section 5.1.2](#).

5.2.2. ALTO Footprint

[I-D.jenkins-alto-cdn-use-cases] describes the use cases for a CDN to be able to obtain network topology and cost information from an ALTO server. These use cases include: "Exposing NSP End User Reachability to a CDN, Exposing CDN End User Reachability to CSPs, CDN deployed

within a Broadband network, CDN delivering Over-The-Top of a NSP's network, and CDN acquiring content from multiple upstream sources". An additional use case may be to advertise CDN End User Reachability to upstream CDNs.

The NSP CDN acting as a dCDN ALTO server may filter and send prefix groups (see [Section 5.1.2](#)) to uCDN ALTO clients according to its policies and with respect to a separate agreement it has with each uCDN. A group may appear as a PID in ALTO network and cost maps.

6. Survey of CDN Delivery Proximity Discovery

6.1. BGP-TE

[I-D.gredler-idr-ls-distribution] proposes a BGP-based mechanism by which link state and traffic engineering information can be collected from networks and shared with external components.

6.2. BGP AIGP

The Accumulated IGP Metric Attribute for BGP [[I-D.ietf-idr-aigp](#)] defines a new TLV attribute in BGP that allows redistribution of IGP costs between ASes belonging to the same managing entity. With AIGP, path selection can take into account IGP costs from other ASes for reaching a certain prefix. For the interconnection of multiple CDNs managed by the same entity ("Inter-Affiliates Interconnection" [[I-D.ietf-cdni-use-cases](#)]), the AIGP information may enable a uCDN to determine which dCDN is topologically the closest to a set of end-users. However, the deployment limitations of AIGP listed in [[I-D.ietf-idr-aigp](#)] ([Section 2.](#) and 3.1.) restrict the applicability of AIGP for this use case.

6.3. BGP Extension for CDNI

See [Section 5.2.1](#).

6.4. ALTO

[I-D.penno-alto-cdn] ([Section 7.1.](#)) describes how ALTO can be used by CDNs in a different administrative domain than the ISP to provide the cost from each CDN node to all known Subscriber PIDs. This mechanism enables the CDN to determine its CDN Delivery Proximity to groups of end users.

[I-D.ietf-alto-deployments] discusses deployment related issues of ALTO for peer-to-peer and CDNs. In [Section 5.1](#), it presents the use of ALTO for a mono-CDN case. The ALTO server may leverage the BGP

information (e.g., BGP community attribute) to group prefixes into PIDs. ALTO cost map permits providing cost information between PIDs and could be used by a dCDN to communicate its CDN Delivery Proximity to an uCDN.

[I-D.seedorf-alto-for-cdni] briefly mentions that ALTO could support selection of downstream CDN but does not indicate the way the ALTO server is fed.

6.5. Generic Capability Advertisement

[I-D.he-cdni-cap-info-advertising] proposes an HTTP/1.1-based protocol which is used to communicate capability information (e.g., resources, footprint, load) "to facilitate selection of the Downstream CDN by the Upstream CDN request routing system".

7. Synthesis

The use cases section shows that in the short term, the need for a Footprint Advertisement Protocol is limited to specific use cases. The survey shows that multiple new proposals addressing CDN Footprint Discovery are being defined, but none is mature and fulfills all the requirements yet.

As a result, we consider that if there is enough interest for developing a CDN Footprint Advertisement Protocol, more work is needed to fulfill the specific requirements that arise in the context of CDNI.

Key building blocks for a Footprint Discovery protocol are clearly identified:

1. Information on the network-level connectivity to groups of prefixes, i.e., to groups of end-users, is required. For instance, BGP-level inter-domain routing data can typically be the source of this type of information, which may be advertised through a protocol based on HTTP, for example, or directly through BGP.
2. A mechanism to group end-users that must be served from the same set of Surrogates (e.g., representing a given CDN) is required. For example, BGP extended community attribute and ALTO PIDs typically provide such mechanisms. The grouping of end-users can disclose confidential information on the CDN or network organization, therefore, a CDN/NSP will provide the groups' definitions only to its trusted partners.

3. A mechanism to discover generic cost information (uCDN Delivery Proximity) for the delivery from a given set of Surrogates (e.g., a CDN) to a given set of end-users is required. For example, ALTO cost maps, legacy BGP, BGP AIGP, and Previdi's MP-BGP extension typically provide such information, which may be advertised through the aforementioned protocols directly or through another protocol (e.g., HTTP based). To fulfill the topology hiding requirements (identified in [[RFC5693](#)] ([Section 5.5.](#)) and [[I-D.penno-alto-cdn](#)] ([Section 6.1.](#))), the advertised information must not disclose confidential information on the CDN's and underlying networks' topology (i.e., it must not permit to derive a detailed network map).

IETF may design a Footprint Discovery mechanism basing on these building blocks. To increase the chances for this mechanism to be deployed by operators, such a mechanism should give enough control on information advertisement and respect operational requirements such as not being too tightly bound to the network.

[8.](#) IANA Considerations

This memo includes no request to IANA.

[9.](#) Security Considerations

Footprint Discovery exposes information about the internals of CDNs. Therefore, it is subject to confidentiality issues.

[10.](#) Acknowledgments

The authors would like to thank Christian Jacquenet, Yannick Le Louedec, Sophie Nachman-Ghnassia, Iuniana Oprescu, Marcin Pilarski, and Emile Stephan for discussions about early versions of this document.

They also thank the contributors of the EU FP7 OCEAN project for valuable inputs.

[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.

- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", [RFC 4760](#), January 2007.

11.2. Informative References

- [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.gredler-idr-ls-distribution]
Gredler, H., Medved, J., Previdi, S., and A. Farrel, "North-Bound Distribution of Link-State and TE Information using BGP", [draft-gredler-idr-ls-distribution-02](#) (work in progress), July 2012.
- [I-D.he-cdni-cap-info-advertising]
He, X., Dawkins, S., Chen, G., Zhang, Y., and W. Ni, "Capability Information Advertising for CDN Interconnection", [draft-he-cdni-cap-info-advertising-01](#) (work in progress), March 2012.
- [I-D.ietf-alto-deployments]
Stiemerling, M., Kiesel, S., and S. Previdi, "ALTO Deployment Considerations", [draft-ietf-alto-deployments-04](#) (work in progress), March 2012.
- [I-D.ietf-cdni-problem-statement]
Niven-Jenkins, B., Faucheur, F., and N. Bitar, "Content Distribution Network Interconnection (CDNI) Problem Statement", [draft-ietf-cdni-problem-statement-08](#) (work in progress), June 2012.
- [I-D.ietf-cdni-requirements]
Leung, K. and Y. Lee, "Content Distribution Network Interconnection (CDNI) Requirements", [draft-ietf-cdni-requirements-03](#) (work in progress), June 2012.
- [I-D.ietf-cdni-use-cases]
Bertrand, G., Emile, S., Burbridge, T., Eardley, P., Ma, K., and G. Watson, "Use Cases for Content Delivery Network Interconnection", [draft-ietf-cdni-use-cases-10](#) (work in progress), August 2012.
- [I-D.ietf-idr-aigp]
Mohapatra, P., Fernando, R., Rosen, E., and J. Uttaro,

"The Accumulated IGP Metric Attribute for BGP",
[draft-ietf-idr-aigp-08](#) (work in progress), June 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.penno-alto-cdn]

Penno, R., Medved, J., Alimi, R., Yang, R., and S.
Previdi, "ALTO and Content Delivery Networks",
[draft-penno-alto-cdn-03](#) (work in progress), March 2011.

[I-D.previdi-cdni-footprint-advertisement]

Previdi, S., Faucheur, F., Faucheur, F., Medved, J., and
L. Faucheur, "CDNI Footprint Advertisement",
[draft-previdi-cdni-footprint-advertisement-01](#) (work in
progress), March 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.

[RFC5693] Seedorf, J. and E. Burger, "Application-Layer Traffic
Optimization (ALTO) Problem Statement", [RFC 5693](#),
October 2009.

Author's Address

Gilles Bertrand
France Telecom - Orange
38-40 rue du General Leclerc
Issy les Moulineaux, 92130
FR

Phone: +33 1 45 29 89 46

Email: gilles.bertrand@orange.com

