Content Delivery Network Interconnection

Footprint versus Capability information and exchange

Bastiaan Wissingh

Institute of Informatics at University of Amsterdam
Netherlands
TNO
Netherlands

***Abstract*— In the current Internet, more and more providers are developing and deploying their own Content Delivery Networks in order to improve the Quality of Experience and decrease the load within their networks. An interesting scenario would be to interconnect these networks to cooperate between these different Content Delivery Networks. The IETF Content Delivery Network Interconnection Working Group has started defining and standardising the framework for interconnecting two Content Delivery Networks. This paper presents an overview of the CDN Interconnection framework and describes and analyses the terms Footprint and Capabilities information as well as the exchange protocols for such information, the extension to M-BGP for CDNi and ALTO.**

***Keywords*— CDN, Interconnection, Footprint, Capabilities, M-BGP, ALTO**

I. INTRODUCTION

Content Delivery Networks (abbreviated as CDN) are overlay networks designed to deliver content to end users with high availability and performance. These networks tend to accomplish that by deploying a large distributed system of servers among multiple data centres across the Internet.

In todays Internet there is many different commercial software available for implementing a CDN as well as many Content Delivery Service Providers (CDSPs) offering such networks to customers around the world [4].

These many different Content Delivery Service Providers vary from big sized providers able to offer content all around the world to small sized providers able to offer content only within certain countries or regions within countries.

For these smaller Content Delivery Service Providers it can be interesting to cooperate with one another in order to deliver content from their clients to more broad areas. For example when a French Content Delivery Service Provider not only wants to deliver the content for its client to end-users in France but also to end-users in the Netherlands.

In order for two different CDNs to be able to connect to each other and exchange content, information about the properties of those networks need to be exchanged. Standardisation of this information is taking place within the Internet Engineering Task Force Content Delivery Network Interconnection Working Group (further referenced to as IETF CDNi Working Group).

These developments define a framework that focusses on the exchange of metadata between CDNs, the exchange of transaction logs and monitoring information, the exchange of request-routing information, the exchange of policies and capabilities and on content management [8].

Section II of this paper describes the research done in this paper. Section III briefly explores different architectures and technologies used within CDNs followed by section IV with a general description of the current situation on interconnecting these networks. Section V dives into selection criteria on which redirect decisions between CDNs can be made. Section VI gives a general explanation on exchange protocols. Section VII discusses the protocols suggested by the IETF CDNi Working Group while the conclusion is presented in section VIII and section IX makes some suggestions for future research.

II. PROPOSED RESEARCH

As mentioned in the introduction, the CDNi framework is focussed on standardising different aspects of the information exchange between CDNs. Two of those aspects are related to the exchange of information to facilitate the proper redirection of an end-user request to a Downstream CDN. This is referred to as Request-Routing information and/or Footprint and Capabilities information.

The above two aspects are the main focus of this research paper and therefor this research looks into the possibilities of using different protocols for the exchange of information between CDNs about their so called Footprint and Capabilities. Different protocols are compared in order to see which protocol is better usable for exchanging such information between the CDNs.

In order to make such comparisons, first criteria for the footprint information as well as the capabilities information are defined after which different protocols can be evaluated. Based on the proposed research above, the following research questions have been defined:

How can Footprint and Capabilities be defined?

Which proposed method is more suitable for exchanging footprints and capabilities between different CDNs?

A. Approach

To be able to address the above questions, different criteria for footprint and capabilities need to be defined in order to start on a valid comparison of the different methods. These criteria have been defined by evaluating the current discussions within the IETF CDNi Working Group as well as by conducting an interview with Stef van der Ziel, founder and owner of Jet-Stream, one of the market leaders in CDN technologies and intelligence [20]. After these criteria for footprint and capabilities are defined, information was gathered on the methods proposed by the IETF CDNi Working Group as also on the manner in which the proposed protocols function. After that a comparison with regards to the above questions has been made.

B. Scope

A comparison between the different proposed exchange methods should be made within the scope of this project. This comparison should pertain to the questions posed above resulting in a conclusion of a more suitable protocol.

Now that the purpose of the research has been made clear, the next section provides an introduction into CDNs and their benefits, architectures and mechanisms.

III. CDNs

A CDN is a network of computers connected across the Internet to transparently cooperate for delivering content to end-users [4]. As previously mentioned is the purpose of a such a network to deliver content in a reliable and timely fashion by replicating content from the origin server to cache servers around the globe located close to the end-users [1].

Within the area of CDNs different terminologies are used to denote the cache servers within the network that provide content to the end-user. These terms include “edge server”, “cache server”, “replica server” and “surrogate server”. Also a server containing the original copy of the content is referred to as “origin server”.

This section briefly looks into the benefits, the components and architecture as well as the technologies used within such CDNs.

A. Benefits of CDNs

CDNs can be used to provide different kinds of services and functionalities, among which are the storage and management of content, the distribution of content among surrogate servers, the delivery of static, dynamic and/or streaming content as also backup and disaster recovery.

By spreading the servers across the Internet instead of hosting on one location performance can be improved since not all end-users have to access the same location. This spreading also helps dealing with sudden peaks in content requests known as flash crowds by giving end-users the possibility to obtain the requested content from the nearest located server.

CDNs can reduce the bandwidth consumption within a network and improve the reliability by making use of a combination of caching and replication through the surrogate servers within the network.

Caching means that a copy of the original data is stored closer to the end-user allowing them to have faster access to the copy of the data. This copy of the data can be stored for example locally on the end-user’s machine or for example on a caching server that is close to the end-user. Replication on the other hand, is a complementary mechanism that manages the copies of the original data throughout a network, so to make sure that all copies are exactly the same for example.

B. Components and Architecture

CDNs generally distinguish three roles, a Content Provider, a Content Delivery Network Provider and the End-Users as shown in figure 1 [4][5].

The Content Provider is the entity responsible for and owner of the content and delivers this content via a so called origin server to the CDN Provider. CDN Providers are organisations providing the infrastructure in order to deliver the content in a timely and reliable fashion to the End-Users who in turn consume the content of the Content Provider.

Research shows that the networks of the CDN Providers consists of four different components, a set of surrogate servers, a request routing infrastructure, a distribution infrastructure as also an accounting infrastructure [4].

The surrogate servers are responsible for delivering copies of the content from the origin server to the end-users. The request routing infrastructure makes sure that the end-users are redirected to the right surrogate server. The distribution infrastructure is responsible for copying the content from the origin server to all the necessary surrogate servers whereas the accounting infrastructure is responsible for the logging, reporting and billing within the network.

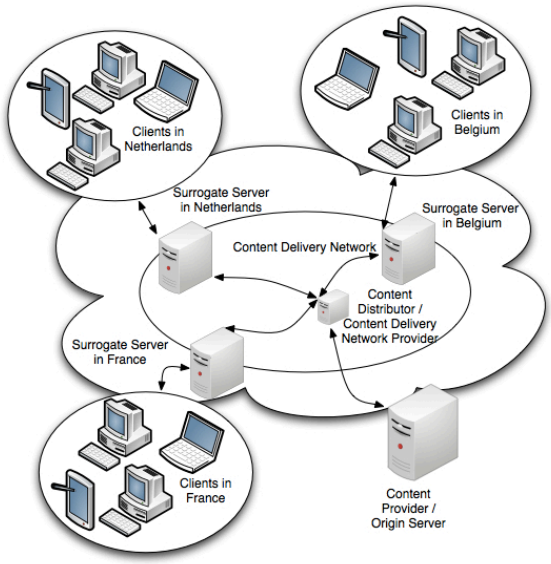

Fig. 1 Illustration of a CDN

According to [4] the organisational structure a CDNs can either be an overlay based structure or a network based

structure. In the overlay based structure the content distribution is handled by the application specific servers and caches whereas in the network based structure this is handled by the network components.

C. Mechanisms in general

CDNs make use of many different mechanisms within the network in order to provide the previously described functionalities [5]. Although specific policies and algorithms used within current CDNs are often proprietary and therefore not publicly available, the general mechanisms can be divided into several categories.

Namely, mechanisms related to placing surrogate servers, mechanisms for updating content throughout the network, mechanisms for actively measuring the network, mechanisms for selecting surrogate servers for handling requests and mechanisms for re-routing end-user requests to the surrogate servers.

For the purpose of this research only the general mechanisms used for redirecting requests of end-users to surrogate servers close to the end-user are discussed. This redirection process basically exists of two parts, deciding on the best surrogate server and redirecting the end-user to the best surrogate server.

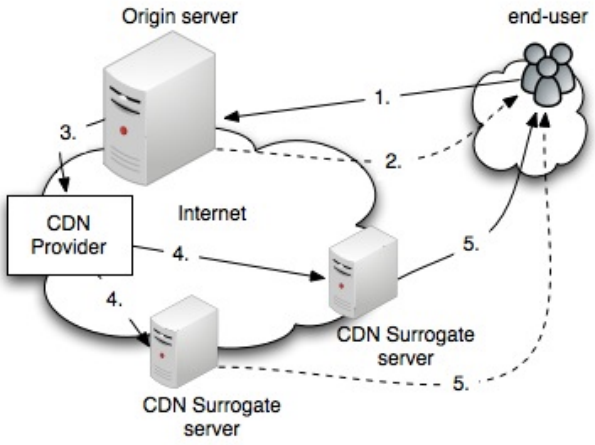

Fig. 2 Basic Request Routing process within a CDN [4][5].

Figure 2 shows the basic steps regarding the process of routing an end-user to the appropriate surrogate server. (1). A client makes a request for content, e.g.. a certain website or media file. (2). The origin server can optionally return (part) of the content based on whether it makes use of full-site or partial-site selection and delivery, explained further on in this section. (3). The origin server examines the request and forwards the request to its CDN Provider. (4). The CDN Provider examines the request to decide on which surrogate complies best with the request and redirects end-user to that surrogate server. (5). The selected surrogate server delivers the content requested by the end-user.

Making a decision on which server becomes the best surrogate server can be based on many different criteria. However, as mentioned before, the selection mechanisms are mostly proprietary and therefor not publicly available. Akamai for example only mentions that its system redirects end-user request to the nearest available server that is suspected to have the requested content based on three parameters, “nearest”, “available” and “likely”. Where “nearest” examines the network topology and dynamic link characteristics, “available” examines the load and network bandwidth and “likely” examines which servers carry the content [7].

The mechanisms to redirect the end-users to the proper surrogate server on the other hand are generally based on different publicly available mechanisms and can be used for both full-site as well as partial-site selection and delivery. Partial-site means that the origin server only hosts static content while dynamic and streaming content is hosted by a CDN. In that case an end-user receives the static content from the origin server while the dynamic and/or streaming content is received through a CDN. In full-site selection and delivery on the other hand, a copy of all content from the origin server is hosted by a CDN, instead of only the dynamic and/or streaming content.

D. End-user redirection mechanisms

Recent research specifies six different mechanisms commonly used for redirecting end-user requests, Global Server Load Balancing, Domain Name System (DNS) based request routing, Hypertext Transport Protocol (HTTP) redirection, Uniform Resource Locator (URL) rewriting and Any-casting mechanism which are discussed briefly [4][1][5].

The first mechanism, Global Server Load Balancing, is based on a master/slave architecture and is used for both deciding on the best surrogate server as well as for redirecting the end-user to that surrogate server. The master node maintains status information like responsiveness and such from all slave nodes positioned in different locations which makes this mechanism global aware. Based on the location of the end-user and the status information of the slave nodes, the master node make a decision on the nearest slave node to which the request is redirected via DNS [6][4].

The DNS-based request routing mechanism in contrary only redirects end-users to an appropriate surrogate server based on the Domain Name given by a separate process that decides on the best surrogate server. So in comparison with Global Server Load Balancing, within the DNS-based redirection there is no dynamic part involved. Based on the manner in which the DNS server is modified, the DNS-Reply can contain one or multiple IP addresses of surrogate servers of which the end-user can contact one to retrieve the content.

The HTTP redirection mechanism makes use of the “HTTP 302 Found” status code, to specify a URL of the surrogate server from which the end-user can receive the requested content.

The URL rewriting redirection mechanism redirects the end-user requests by modifying the URLs embedded in content items so that they redirect to surrogate servers of

which the end-users can gain the content items. This can be either pro-active (passive) manner by which the embedded URLs of the main item are modified when the content is placed on the origin server or reactive (dynamic) manner by which the embedded URLs are modified upon a end-user request.

The Any-cast redirection mechanism is an Internet Protocol Layer approach to sending data or requests to a single entity from a group of entities. This can be done by assigning a group of entities with the same IP address so that the entity closest to the sending entity with receive the data or request and can setup a connection with the sending entity. Related to the Surrogate servers, the closest server can provide the content to the sending entity.

Now that an introduction has been given into CDNs and their benefits, architectures and mechanisms the next section describes the framework currently being defined by the IETF CDNi Working Group for standardising the communication between interconnect CDNs.

IV. INTRODUCTION TO CDN INTERCONNECT

Within the current standardisation process run by the IETF CDNi Working Group, multiple companies from different industries are involved. Together these companies are working on standardising a framework for interconnecting CDNs. This section discusses the problem statement for interconnecting CDNs as well as the framework for the interconnection as suggested and described by the IETF CDNi Working Group.

A. Problem statement

One could think of multiple situations in which it could be desirable for a CDN to be able to interconnect with another CDN in order to deliver content from its Content Provider to end-users. An example of such a situation has already been mentioned in the introduction.

This section provides a short summary of different situations in which an CDN Interconnection could be preferable, as also discussed within the IETF CDNi Working Group [26][27][28].

A CDN Provider can decide to interconnect his network with another CDN (Provider) for example in case of a disaster or flash crowd within his own network. In that case the other network can take over his service so that end-users are still able to retrieve the content from the Content Service Provider. Another situation could be where a Service Provider operates over multiple geographical locations and wants to interconnect those Network Service Providers in order to seamlessly deliver its content to end-users hopping from one location to another.

Within the IETF CDNi Working Group the situations for having an CDN Interconnection have been categorised into three categories related to Footprint extensions, Offload situations and Capability extensions which are described next [27].

The term Footprint is defined as geographical coverage of a CDN and is further discussed later in this paper. This category therefor describes situations regarding CDN Providers that are able to expand their services beyond their own coverage by interconnecting with other CDN Providers as also situations which can benefit Internet Service Providers by letting them reduce traffic load within their network and influence and control traffic by interconnecting with other Internet Service Providers.

The term Offload is defined as handing over to another entity, therefor this category describes situations in which CDNs can for example increase capacity during traffic peaks (flash crowds), internal failures or in specific regions by interconnecting with another CDN.

The term Capability is defined as functionalities a CDN is able to provide. This category therefor describes situations in which the CDN is able to expand its capabilities by interconnecting with other CDNs. Examples include expanding support of devices and technologies like able to deliver streaming “MP4” content to Apple’s iOS devices, expanding Quality of Experience and Quality of Service.

B. Proposed framework

To be able to address the interconnection between CDNs in situations as described in the previous section, the IETF CDNi Working Group has proposed an CDN Interconnection Framework which is described in this section [29].

The CDN Interconnection Framework defines four different interfaces needed to interconnect two CDNs with each other. The Control interface, the Logging interface, the Request Routing interface and the CDNI Metadata. These interfaces are used for communication between an Upstream CDN and a Downstream CDN.

An Upstream CDN (further referred to as uCDN) is defined as the network that redirects an end-user request to another CDN while a Downstream CDN (further referred to as dCDN) is defined as the network to which a end-user request is redirected by another CDN [26].

The Control interface of the Framework is responsible for controlling the other CDN components and interact with other CDNs within the interconnection. Via this interface different CDNs initiate an interconnection and bootstrap the other interfaces of the Framework.

The Logging interface of the Framework is responsible for the exchange of log information between the interconnected CDNs. This information is related to the process of delivery of content, general activity within the network and diagnostics.

The Request Routing interface of the Framework is responsible for deciding on which dCDN is best for redirecting the end-user to as well as taking care of redirecting the end-user to that dCDN. Therefor this interface lets the interconnected CDNs communicate information about their Footprint and Capabilities (which is discussed in the next

section) as well as information needed to redirect end-users to specific CDNs and surrogate servers within those networks.

The fourth interface defined within the Framework is the Metadata interface, which is responsible for the exchange of metadata about content between the interconnected CDNs as well as the exchange of content itself.

Three of the interfaces of the Framework can be related to the components of which the networks of CDN Providers exist, as previously described. The Logging interface can be related to the so called accounting infrastructure, responsible for the logging, reporting and billing facilities within the network. The Request Routing interface can be related to the request routing infrastructure, responsible for redirecting the end-users to the right surrogate server and the CDNI Metadata interface can be related to the distribution infrastructure, responsible for replicating the content from the origin server to the surrogate servers.

Now that the general structure of the framework has been described, the next section looks into the criteria on which an uCDN should be able to select a dCDN to redirect an end-user request to.

V. DOWNSTREAM CDN SELECTION CRITERIA

As described in the previous section, according to the CDN Interconnection framework, should CDNs exchange footprint information as well as information about their capabilities in order to be able to make a granular selection of a dCDN.

Within the IETF CDNi Working Group discussion there is a distinction between the definition of the terms Footprint and Capabilities on the one hand and the process of advertising that information on the other hand. The advertising process is referred to as Footprint and Capabilities advertisement [30].

The purpose of that advertising process is to enable the uCDN to decide which dCDN (connected via an interconnection) it wants to redirect a request of an end-user to as also to determine whether that dCDN can handle that redirection from the end-user. In order for the uCDN to make such a decision and determine the possibilities it needs to receive information from the dCDN regarding its capabilities. This information about the capabilities is clearly being divided into so called Footprint information and Capabilities information. The general idea is that an uCDN can make an initial decision for a certain dCDN by looking at the Footprint information while additional Capabilities information can be used when the Footprint information is insufficient to make a delegation decision [30][29]. This clear distinction made by the IETF CDNi Working Group, is further discussed at the end of this section.

This section looks into the requirements specified by the IETF CDNi Working Group for the CDNi framework which are related to footprint and capabilities. Whereafter the terms footprint and capabilities are discussed to try to argue on how these terms could best be defined and whether the clear distinction between the two is valid.

A. Requirements of the framework

As mentioned before has the IETF CDNi Working Group specified a certain set of requirements that need to be fulfilled by the CDNi framework in general as well as by the specified interfaces within the framework. This section gives a general description of the requirements that are related to the terms Footprint, Capability and the exchange of both [32].

An important general requirement of the framework to note is that the framework shall not require intra CDN information like the topology of surrogate servers, the status of surrogate servers and such to be communicated or exposed to other CDNs in order to provide efficient delivery of content. This is an interesting requirement since one of the definitions of Footprint information (as later discussed) is the set of IP information from surrogate servers within the dCDN.

Besides the general requirements there are also requirements defined for each specific interface within the framework. One of those specific requirements defined for the Request Routing interface specifies that the Request Routing interface must allow the dCDN to communicate information to the uCDN about its ability to handle requests as also information to facilitate the selection process of the uCDN. So the request routing interface must exchange the footprint and capabilities information between two CDNs.

Also one of the requirements of the Metadata Distribution interface related to Footprint and Capabilities information is that the interface shall make it possible for the Upstream and dCDN to exchange information about content distribution policies like geo-blocking information, availability windows and delegation white- and blacklists.

Besides the previous requirements it is required for the whole framework that it should support secure operation over unsecured IP connectivity via mechanisms like authentication, confidentiality, integrity and spoofing and reply protection. It is however not specified how the different interfaces within the framework should provide support for it.

B. Footprint

The term Footprint is by the IETF CDNi Working Group generally defined as geographic region for which a CDN is able to deliver content (either directly or via delegation to another CDN). The term geographic region however is somewhat vague, as there are different definitions which all somehow boil down to the statement that a geographic region is a certain delimited area of the earth. The CDNi Working Group suggests that this area information can be covered by either a set of country, state and city code combinations (ISO 3166-2), a set of Autonomous System numbers [33] or a set of IP subnets [31][32][29].

When one compares “certain delimited area of the earth” to the three previous given suggestions of representation of geographic region, only the first example (a set of country, state and city code combinations) matches. The other two examples of representation can not be reflected to a delimited area of the earth as is explained below.

Let us first look at the suggestion of a set of country, state and city code combinations as Footprint. The different country and state code combinations from all around the world are defined by the International Organisation of Standardisation in the ISO 3166 standard. This standard specifies codes for the representation of names of countries and their subdivisions and consists of two parts. The first part describes the Country codes while the second part describes the Country subdivision codes [34].

It is however interesting to note that the second part of the standard is not consistently defined. For some countries the subdivision is based on provinces while for others it is based on regions, departments, districts or even a combination of the previous. Due to this inconsistent definition a combination of country and state codes can not be used as entity to base the footprint information of a CDN on. It is for example not possible to compare footprints of different CDNs in such case.

The second suggestion is to make use of so called Autonomous System numbers as Footprint information. Autonomous Systems numbers are numbers used within the Internet to indicate a collection of IP prefixes which are under the control of one or more network operators and have a single and clearly defined routing policy towards the Internet. All Autonomous System numbers within the Internet have a globally unique number that is used to identify the Autonomous System as well as for exchanging routing information [33].

In comparison with a set of country, state and city combinations, Autonomous Systems also represent a certain region however not necessarily geographic. Autonomous System numbers look at regions from a more Internet perspective instead of from a human perspective. The numbers indeed represent different regions, however this regions are not topologically related. By Internet perspective is meant for example that when you have an internet connection at home via Internet Provider X (having AS Number 1) and you would like to connect to your neighbour who lives next door and has an internet connection provided Internet Provider Y (having AS Number 9), it could be that although your neighbour is geographically close by, your connection has to pass different networks in order to reach your neighbour.

Since Autonomous System numbers indicate a collection of IP prefixes, it is possible to make a mapping between these numbers and the corresponding IP prefixes and between those IP prefixes and the IP information of an end-user in order to select an Autonomous System than can reach the IP address of the end-user. Therefor Autonomous System numbers are a good candidate to be used as the so called Footprint information.

The third suggestion for Footprint information is to make use of a set of IP subnets. These sets can be specified as full IP addresses or prefixes either IPv4 or IPv6 which indicates for example end-user requests a dCDN is able to serve or IP addresses of the surrogate servers that are deployed within the dCDN. When the set of IP subnets represent the surrogate servers deployed within the dCDN however, the information is in contradiction with one of the requirements of the framework as described before [29]. The framework shall not require intra CDN information like to topology of surrogate servers, the status of surrogate servers and such to be communicated or exposed to other CDNs in order to provide efficient delivery of content.

If one compares this suggestion to the suggestion to make use of Autonomous System numbers, the information exchanged in this third suggestion is much more specific. However providing such specific information via sets of IP subnets does not mean that an uCDN can make a better decision than when the Footprint information is provided in the form of Autonomous System numbers. These sets of IP subnets for example do not state anything about whether those IP subnets are directly connected to a certain end-user or whether there is an certain amount of hops in between. It does however generates a lot more information that needs to be exchanged between CDNs. This makes it a less suitable candidate for the expression of Footprint information within the CDN Interconnection framework.

Now that all three suggested categories have been described and analysed, it is interesting to notice that there is such a difference in granularity between the three. Changing from very general information of a combination between country and states to very detailed information of lists of IP prefixes. As discussed the best candidate for the Footprint information seems to be the Autonomous System numbers, since they provide a consistent devision of regions with sufficient information to be able to make a mapping between the location (IP address) of the requesting end-user and the region that is able to serve that end-user.

Since the definition of the term geographic is vague and does not provide sufficient information from an Internet perspective, as discussed, it is not advised to be used as entity to base the footprint of a CDN on. So a better definition of Footprint would be region instead of geographic region for which a CDN is able to deliver content (either directly or via delegation to another CDN).

Although at this point the Autonomous System numbers seem to provide the better way on how to provide Footprint information, a question remains whether a dCDN should only provide information on the Autonomous Systems it is directly connected to or also on the Autonomous Systems that it can delegate requests to. As long as there is a mechanism in place which provides the possibility to indicate of which Autonomous Systems the dCDN itself is part of and which Autonomous Systems it is able to redirect his request to information on both should be provided. This provides an uCDN with the possibility of a much broader coverage without having to setup an interconnection agreement with another dCDN.

C. Capabilities

As already mentioned is the idea of exchanging information about the capabilities of a dCDN with an uCDN to make it possible for the uCDN to make a more granulated decision on the dCDN to which it will redirect an end-users request. As the term suggests are capabilities related to the features, services and states a CDN can or cannot meet. The IETF CDNi Working Group has suggested to split this information into four different categories, information about the caches, the resources, the network and the administrative capabilities of the dCDN as described next [30].

Information about the caches is being referred to as information about the load, the available resources in terms of storage and failure conditions. Resource information on the other hand, is being referred to as information about playback devices, delivery technologies and content types the dCDN is able to support such as the ability to provide streaming "MP4" content to Apple's iOS devices.

Network information is information about a certain quality of service, distribution and delivery priorities and the streaming bandwidth that is supported. Whereas administrative capabilities, is being referred to as policies and administrative limits (such as the maximum volume of aggregated content the dCDN is able to server) and variables such as fees.

These suggested categories comply with the requirements specified for the CDNi Framework. As mentioned in the first section of this section, the capabilities information should allow the dCDN to communicate information about its ability to hand delegated requests by communication information about its current status such as load as also to communicate information to facilitate the selection process by the uCDN by communicating supported content types, metrics, affinities and policies.

D. Distinction between Footprint and Capabilities

Now that we know what the definition of Footprint and Capabilities is and what the purpose of their use is, we could question whether the clear distinction in the purpose of both entities is valid. As mentioned in the introduction of the Footprint and Capabilities section the general idea is that an uCDN can make an initial decision for a certain dCDN by looking at the Footprint information while additional Capabilities information can be used when the Footprint information is insufficient to make a delegation decision

If one however looks at these purposes from a more practical aspect, the question could be posed whether it is useful to try to make a selection based only on the Footprint information and if that information is insufficient also compare capabilities of dCDNs. The reason for that is that there are sufficient examples available of situations in which a selection based on Footprint information would lead to an sub-optimal or incorrect decision. Some of these examples have been discussed with Stef van der Ziel founder and owner of Jet-Stream, one of the market leaders in CDN technologies and intelligence [42]. This meeting was intended to discuss, among other subjects, the ideas of footprint and capabilities information from a more practical perspective.

One of the examples that came forth during this meeting is when one considers CDN A with a footprint in Belgium that provides delivery of static content to its end-users and CDN B with a footprint in France that provides both the delivery of static content as well as streaming content to its end-users. If the uCDN is looking for a dCDN that can deliver streaming content to and end-user in Belgium, it would select the Belgium provider if the selection is initially based on the Footprint. There is no footprint overlap, so the footprint information seems sufficient at first sight. This decision however would be incorrect since that CDN is not able to provide the required content to the end-user.

A better approach would be not to divide the selection process into these two concepts, but to make the footprint information part of the capabilities requirements. Then the selection of the dCDN would not only be based on the footprint information but will be based on a selection of capabilities. In this case the selection could be much more sophisticated and situations as given in previous examples could more easily be avoided.

Now that the terms Footprint and Capabilities have been discussed in this section, the next section provides a general description of the protocols Border Gateway Protocol and Application Layer Traffic Optimisation. The purpose is to provide a basic understanding of the protocols so that section "*VII. Suggested exchange protocols"* can discuss the protocols suggested by the IETF CDNi Working Group in depth.

VI. EXCHANGE PROTOCOLS IN GENERAL

Within the IETF CDNi Working Group as set of protocols has been suggested and designed to facilitate this exchange of Footprint and Capabilities information. The suggested protocols are related to and variations on the Border Gateway Protocol version 4 [9] and the Application Layer Traffic Optimisation protocol [12]. Therefor this section first explains the Border Gateway protocol whereafter the Application Layer Traffic Optimisation protocol is explained. The suggested protocols by the IETF CDNi Working Group are discussed in section "*VII. Suggested Exchange protocols"*.

A. Border Gateway Protocol version 4

As of 2006 the fourth version of the Border Gateway Protocol (abbreviated as BGP) has been standardised by the IETF. The purpose of BGP is to provide connectivity between different networks of providers so that these networks are able to connect to the Internet. In order to provide this, BGP exchanges routing information between these providers across the Internet in a decentralised way [23].

The different providers are seen by BGP as independent Autonomous Systems. As mentioned before, Autonomous System collections of IP prefixes which are under the control of one or more network operators and have a single and clearly defined routing policy towards the Internet. All Autonomous System numbers within the Internet have a

globally unique number that is used to identify the Autonomous System as well as for exchanging routing information [33].

To exchange the routing information between these Autonomous Systems, BGP makes use of a so called Path Vector Routing algorithm. As the name suggests does a path vector maintain path information in order to make a routing decision. One of the ideas behind a path vector algorithm is that it should be able to prevent loops since the whole path is known which makes the algorithm capable of detecting duplicate items within a path.

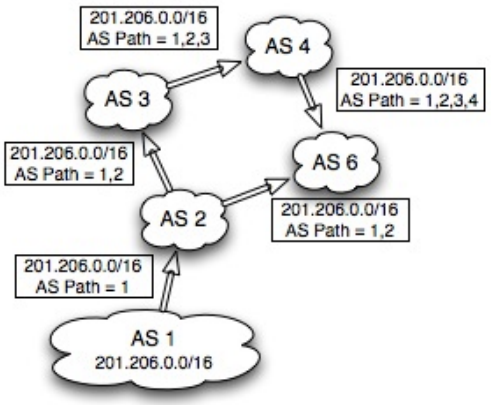

Fig. 3. Basic operation of Path Vector principle within BGP

BGP can be run in two ways, between different Autonomous Systems which is being referred to as external BGP (eBGP) or within Autonomous Systems which is being referred to as internal BGP (iBGP). eBGP is used to exchange prefixes and implement policies between different Autonomous Systems, as also to prevent routing loops by making use of a path vector algorithm to detect duplicates in the paths between Autonomous Systems. iBGP on the other hand is used to distribute prefixes learned from eBGP within an Autonomous System. Since BGP makes use of Autonomous System paths to detect loops within the network, the loop detection process only works with eBGP as Autonomous System numbers are not communicated into other Autonomous Systems. In order to prevent loops within iBGP, it is therefor necessary that all BGP routers within and iBGP network are connected in a full-mesh (meaning every router is connected to every other router).

Different BGP Peers can exchange multiple types of messages, for further purpose of this paper however only the the OPEN and UPDATE messages are discussed. The BGP OPEN message is used between two BGP neighbours to setup a peering relationship. With this message the BGP Neighbours exchange, among other things, the presence of the so called Optional parameters to negotiate on additional functionalities. This is for example used to check whether the other BGP peer understands Multiprotocol BGP (as described in the next section).

The original BGP standard states that when one of the optional parameters is unrecognised by one of the BGP peers, the connection should be terminated [23]. This way it would not be possible for example for a BGP router that uses Multiprotocol BGP to communicate with a BGP router that does not understand Multiprotocol BGP. To solve such issues, an optional parameter called Capabilities has been defined which provides the introduction of new capabilities in BGP by providing graceful capability advertisement without requiring that BGP peering be terminated [43].

Between different BGP peers, BGP routers directly connected to each other, so called Network Layer Reachability Information (abbreviated as NLRI) is exchanged along with Path attributes that specify additional information about the information contained by the NLRI. The NLRI is information about IPv4 prefixes that the BGP router can or can no longer reach [23].

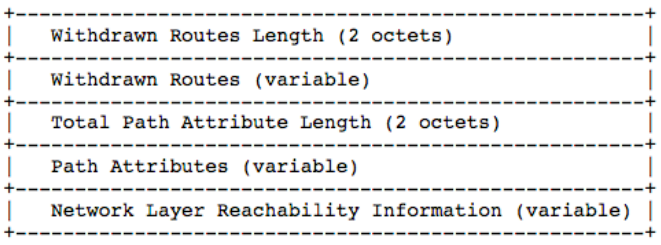

Fig. 4. Layout of the payload of BGP Update message [23].

The path attributes are divided into four different categories of attributes, well-known mandatory, well-known discretionary, optional transitive and optional non-transitive. An attribute categorised as well-known must be recognised by all BGP implementations and when updated by a BGP peer send to all its BGP neighbours. The well-known category can be divided into mandatory, meaning that the attribute must be included in every update message, or discretionary, meaning that the attribute may or may not be sent in an update message. In contrast to the well-known category, there is the optional category. When an attribute falls into this category it may be part of an update message and is not required or expected to be supported by all implementations of BGP.

An optional transitive attribute should be accepted either when recognised or not recognised, however this is not mandatory and depends on the implementation of BGP. An optional non-transitive attribute should however be ignored and not communicated to other BGP neighbours when not recognised by a BGP peer. Now that we have a general view of the working of BGP, lets look at an important extension on this protocol.

B. Multiprotocol extension for BGP

As of 2007 the IETF has standardised extensions to BGP-4 in order to provide support for multiple network-layer protocols instead of only IPv4. These extensions are being referred to as Multiprotocol extensions for BGP-4 (abbreviated as MBGP) [25].

This standard defined two new attributes for use within the BGP-4 messages, namely the Multiprotocol Reachable NLRI attribute (abbreviated as MP_REACH_NLRI) and the Multiprotocol Unreachable NLRI attribute (abbreviated as MP_UNREACH_NLRI). The MP_REACH_NLRI attribute provides information about sets of reachable destinations together with the next hop while the MP_UNREACH_NLRI provides information about a set of unreachable destinations.

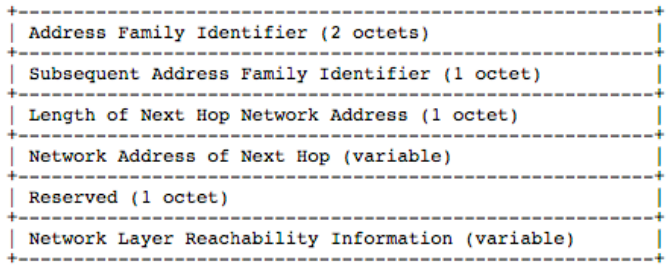

Fig. 5. Layout of MP_REACH_NLRI [25].

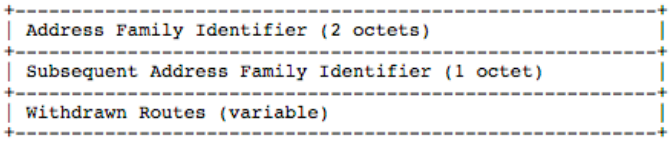

Fig. 6. Layout of MP_UNREACH_NLRI [25].

Both of these attributes are part of the UPDATE message exchanged between different BGP peers. As described in the previous section, in order for two BGP peers to exchange UPDATE messages with the Multiprotocol extension, the peers negotiate the presence of the extension on both sides via the Capabilities advertisement [43].

Both the MP_REACH_NLRI and MP_UNREACH_NLRI attributes are categorised as optional non-transitive which, as described in the previous section, means that when a BGP peer does not recognise the attributes is must ignore them and not distribute them to other BGP peers.

C. Application Layer Traffic Optimisation

Within the IETF there is an active Working Group focussing on the development and standardisation of the so called Application Layer Traffic Optimisation protocol (abbreviated as ALTO). The intention of this protocol is to provide an information sharing service that makes applications capable of performing a "better-than-random" selection of peers [12][18].

In general applications do not have reliable information of the underlaying network which forces them to select peers randomly or based on partial observations which can result in suboptimal choices. "Better-than-random" is being referred to as the opposite of this situation.

The ALTO protocol basically consists of two parts, a discovery mechanism for applications to find a reliable information source (referred to as ALTO server) and a protocol to query such information sources for information that can facilitate in making a better-than-random selection of peers. This information that can be provided by ALTO servers is related to operator policies, geographical location, network proximity as also transmission costs. The ALTO server can retrieve this information from entities like network operators and third parties [12].

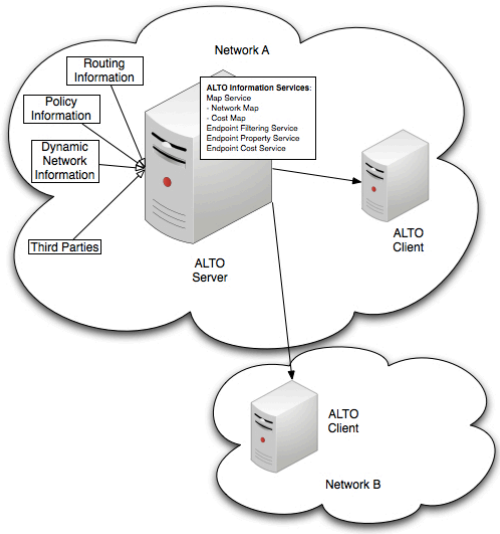

Fig. 7. Overview of general ALTO Architecture

As can be seen in figure 7, the ALTO protocol has a client-server architecture. The communication between these clients and servers is based on HTTP. The protocol provides four different information services, a Map service, a Map filtering service, an Endpoint property service and an Endpoint cost service.

The Map service which is seen as the ALTO-Core exists of two parts, a Network map and a Cost map. The Network map provides information on a full set of network locations while the Cost map provides information on the costs between the different entities in the Network map. The Map filtering service allows the clients to query the server for the Network and Cost maps. The Endpoint property service allows clients to query the server for properties of individual endpoints stored in the Network map and the Endpoint cost service allows clients to query the server for costs between specific endpoints.

The Network map consists of multiple so called provider identified network location identifiers (PIDs) which contains a grouped set of endpoint addresses. Based on this Network map, the Cost map is filled with information about the path costs between the different PIDs.

Based on the previous described structure of the protocol and kind of information stored by the ALTO servers, it should be possible for an ALTO client to perform "better-than-random" selection of peers.

VII. Exchange protocols suggested by IETF CDNi Working Group

Based on the definition of footprint and capabilities defined in section "*V. Downstream CDN selection criteria"* and a general description of BGP and ALTO in section "*VI. Exchange protocols in general"*, this section evaluates the specific protocols suggested by the IETF CDNi Working Group and try to argue which protocol can be best used with regards to the exchange of the previously described information.

An interesting observation beforehand is that the CDNi framework suggested by the IETF CDNi Working Group only considers redirect mechanisms that are based on DNS and HTTP [29]. Although multiple mechanisms exist, as described in section three, the mechanisms based on DNS and HTTP are the most commonly used [4]. Also for the evaluation of the different suggested protocols, a CDN is considered as an application-layer network on top of the Internet [21].

Within the IETF CDNi Working Group as set of protocols have been suggested and designed to facilitate the exchange of the Footprint and Capabilities information between different CDNs. The protocols can be divided into three categories, protocols supporting the exchange of footprint information, protocols supporting the exchange of capabilities information and protocols that support the exchange of both, see table below.

Protocol	Footprint Exchange	Capabilities Exchange
Standard BGP	X	
BGP Extended Communities Attribute	X	
BGP-TE	X	
BGP-AIGP	X	
HTTP		X
Extension to M-BGP for CDNi	X	X
ALTO	X	X

Table 1. Exchange protocols suggested by IETF CDNi Working Group

This section first describes the protocols capable of exchanging only Footprint or Capabilities information whereafter the two protocols capable of exchanging both Footprint as well as Capabilities information are described in more depth. The end of this section then argues on which of the latter two protocols lends itself best for the exchange of both Footprint as well as Capabilities information.

A. Proposed Footprint exchange protocols

The first category covers the protocols suggested by the IETF CDNi Working Group that can only be used to exchange information about the footprint of CDNs as also delivery proximity information regarding these footprints. The suggested protocols are based on standard BGP with additional attributes, the Extended Communities Attribute, the Traffic Engineering Attribute and the Accumulated IGP Metric Attribute [35].

Standard BGP could be used as protocol to exchange both footprint information as well as delivery proximity information to that footprints. A dCDN that covers a certain Autonomous System can advertise that information via BGP to an uCDN. This makes it possible for the uCDN to map the Autonomous System number to IP prefixes which can be matched to the IP information of the end-user. As BGP also advertises the Autonomous System paths, it is also possible to determine the delivery proximity to certain Autonomous Systems from the uCDN perspective.

Another suggested protocol to exchange only the footprint information is to make use of the additional BGP Extended Communities Attribute which provides a mechanism that is capable of labelling information contained in an BGP packet [24]. This provides a CDN with the possibility to aggregate information on prefixes within the Autonomous System and communicate the aggregated information instead of more detailed information per prefix.

The BGP Traffic Engineering Attribute provides the ability for BGP to carry and make use of traffic engineering information like minimum and maximum bandwidths and priorities. The suggestion is that this information can be used to collect link state and traffic engineering from the internal networks and share that with external components so that delivery proximity information can be provided [10].

Another suggestion to exchange delivery proximity information is the extension on BGP that makes use of the Accumulated IGP Metric Attribute. This attribute allows Internal Gateway Protocol costs to be exchanged between Autonomous Systems that belong to the same managing entity. This way an uCDN can take IGP costs of other Autonomous Systems into account during the selection of a dCDN [11].

B. Proposed Capabilities exchange protocols

Besides protocols that are only capable of exchanging Footprint information between CDNs, the IETF CDNi Working Group also suggest a protocol that can be used in report and query mode for the exchange of capabilities information. The candidate proposed by the Working Group is HTTP [31].

The reason the Working Group suggests HTTP is that the CDN capability information is related to a specific application, namely the CDNi. Therefor it should be exchanged via an application layer protocol rather than an

underlying protocol which decouples the information from the application.

For the report mode of the protocol, one could use the HTTP POST method. This way a dCDN can for example advertise his capabilities to the uCDN. For the query mode, so that an uCDN can request capability information from the dCDN, the HTTP GET method could be used [31].

C. Footprint and Capabilities exchange

Besides the previous described protocols that are either capable of exchanging footprint information or capabilities information, the IETF CDNi Working Group also suggests two protocols that can be used for both exchanging footprint information as well as capabilities information. The first suggestion is an extension on BGP for CDNi while the second suggestion is to make use of the ALTO protocol.

1) Multiprotocol extension for BGP

The extension to BGP for the CDN interconnection is based on the usage of the previously described Multiprotocol extension for BGP and referred to as a CDN level complement to the network level (standard) BGP [38].

It defines two new subsequent address family identifiers for IPv4 and IPv6 CDNi purposes and a new set of NLRI that contains either Footprint Element (FPE) information, Footprint Reachability (FPR) information or CDN Capabilities (CAP) information. These NLRI entities are distinguished as FPE-NLRI, FPR-NLRI and CAP-NLRI whereby the FPE-NLRI in this case contains an arbitrary set of prefixes that is part of the CDN Footprint, the FPR-NLRI indicates the way in which a CDN can reach one or more prefixes of the Footprint and the CAP contains a set of capabilities supported by the dCDN.

The above mentioned information is exchange via three different Multiprotocol BGP messages, the Footprint Element Advertisement, the Footprint Reachability Advertisement and the Capabilities Advertisement. These three messages are all Multi-protocol BGP messages containing the different NLRI sections as described above. CDNs participating in the exchange of this information via BGP maintain a different database for each of the advertisement types.

Figure 8 shows a graphical representation of M-BGP within CDNi based on an example given in [38]. Hereby are dCDN2, dCDN3 and uCDN connected via CDNi. dCDN2 has a direct connection with dCDN3 and dCDN3 has also a direct connection with uCDN.

Also AS100 and AS400 do not have a CDN and AS100 advertises all his prefixes via BGP to dCDN3 and AS400 while it only advertises the prefixes "1.1.1.0/24 and 3.3.3.0/24" to dCDN2. The figure on the left shows the three different databases of each CDN after they exchanged FPE, FPR and CAP information.

Fig. 8. Example of M-BGP within interconnected CDNs

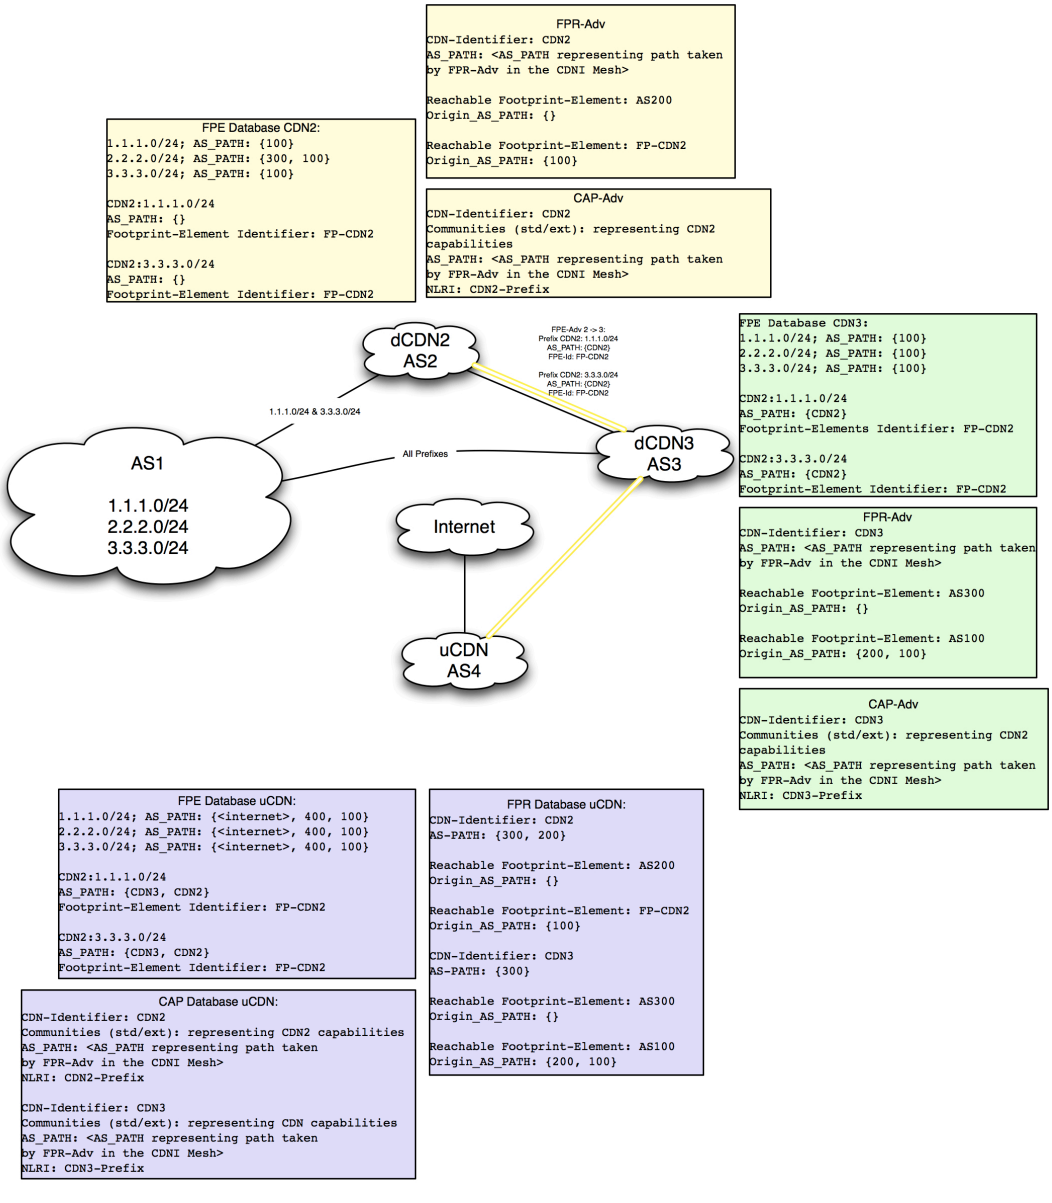

2) *Application Layer Traffic Optimisation*

Now that we have a general idea of how the extension to BGP for CDNi works we can look at the ALTO protocol in relation to CDNi. The general idea of ALTO in relation to CDNi is that the uCDN acts as a client of different dCDNs [39][40].

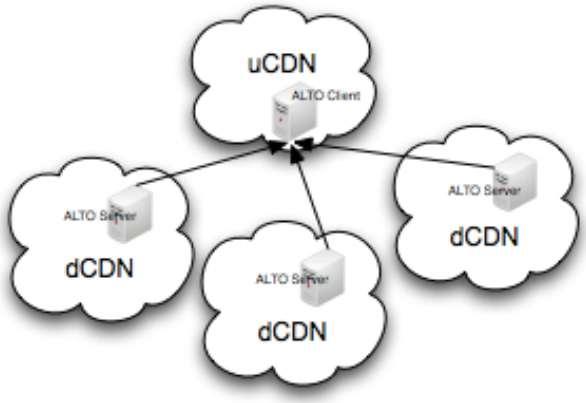

Fig. 9. General ALTO CDNi architecture

The uCDN can in this way gather network and cost maps from the dCDNs providing it with information on the coverage and costs to deliver content via a certain dCDN. Based on this information it should be possible for the uCDN to select the better dCDN to redirect the end-user request to.

Although the working of ALTO for CDNi does not really differ from the general working of ALTO, there are some considerations to be made when using ALTO for Footprint and Capabilities exchange between CDNs as described in [40]. Two important categories of these considerations are related to the ALTO client-server session setup as well as to scalability of the protocol.

In order for two CDNs to make use of the ALTO protocol to exchange information on footprints and capabilities a session between an ALTO server and client needs to be setup. For such a session to be setup, there needs to be an agreement on the configuration of the session, the information that can be exchanged as also on how this information can be exchanged [40]. The current design of the ALTO protocol does not specify these options, [13] but the IETF CDNi Working Group proposes that the options configured in a session between the uCDN client and the dCDN server should somehow reflect the agreements made between these CDNs.

Also since the uCDN is client of multiple dCDNs (as shown in figure 9), this could lead to the uCDN receiving lots of information from the dCDN server that may even result in failures of the uCDN client as it is unable to process the amounts of data. Therefor the amount and frequency of exchanged information must be specified for scalability and responsiveness issues.

The IETF CDNi Working Group makes a couple of suggestions on how to specify the amounts and frequencies of exchanged information related to filter settings of PIDs, summaries of PIDs and incremental updates. The current draft of the ALTO protocol specifies that the client can filter out the PIDs of interest by using the Map Filtering Service (as described in sectionVI. C.). This service however needs to receive the parameters on which to filter every time the client queries for specific PID information. To reduce the amount of data however, it could be desirable to let the client specify the parameters on which to filter on forehand instead of every time the client queries the server. In order to provide this functionality, there are two options for having PID filters at the session level: The filters are agreed by uCDN and dCDN operators and set in the configuration of the session or Filters are dynamically configured during the session by an uCDN ALTO client. It requires the creation of a 'PIDs filters Setting' service in the dCDN ALTO server.

Besides the possibility to only receive filtered PID information, it should also be possible for the uCDN to receive summaries of the Network maps to reduce the amount of exchanged data and let the uCDN decide on the PIDs of which it would like to receive more detailed information. Also the amount of updates from the dCDN to the uCDN should be able to be restricted by the uCDN so that the amount will not exceed the capacity of the uCDN.

D. *Extension to M-BGP for CDNi versus ALTO*

Looking at the information provided in the previous two sections, we are now able to compare the Multiprotocol extension for BGP with the Application Layer Traffic Optimisation Protocol to argue on which of these protocols lends itself best for the exchange of footprint and capabilities information.

As described in the introduction of this paper, for comparison purposes of the previous described protocols a CDN is considered as an overlay network on top of another network such as the Internet [21]. A CDN is therefore considered to be an application which suggests that data between different CDNs should be exchanged via an application layer protocol. Both the M-BGP and the ALTO protocol are considered to be application layer protocols as defined in the Internet Protocol Suite, which defines a set of protocols used for communication on network such as the Internet.

As described in one of the previous sections about the extension to M-BGP for CDNi, there are some drawbacks to using M-BGP for exchanging information about the footprint as well as capabilities between different CDNs. So are the newly defined attributes categorised as optional non-transitive which means that when a BGP peer does not recognise the additional attributes is must ignore them and not distribute them to other BGP peers. This way there is no guarantee that when a dCDN sends information in M-BGP messages to an uCDN that the information is not dropped by one of the BGP routers that resides on the path between the dCDN and the uCDN.

Also since the extension to M-BGP for CDNi is only useful for CDNs, it will most probably not be implemented by other entities than CDN Providers which brings us back to the previous point that there will networks that do not understand M-BGP and therefor not pass on the M-BGP information.

The ALTO protocol on the other hand seems to be a better candidate for exchanging footprint as well as capabilities information. Although this protocol is based on a client-server architecture which could lead to single point of failure situations, the protocol also has some advantages compared to the extension on M-BGP for CDNi.

An example of a single point of failure situation is when the ALTO client receiving to many updates from multiple ALTO servers leading to unresponsiveness of the ALTO client. This has been further described in section VII.

An advantage of ALTO is that the protocol is at the time of writing still in the phase of definition and standardisation. This makes it more easy to suggest additional functionalities that can be used within the CDNi as the protocol is still under discussion. Also when compared to the deployment situation of M-BGP, there is no downside of only implementing the ALTO protocol within different CDNs as the protocol does not rely on other entities that should understand the protocol except for the client and the server.

Also the framework of the ALTO protocol is defined in a more flexible way to support future extensions. So unless there is an already standardised protocol that can be used for the exchange of footprint and capabilities information (see section "*IX. Future Research"*) the current conclusion is that from the extension to M-BGP for CDNi and ALTO, the latter one is more suitable for exchanging this type of information between CDNs.

VIII. CONCLUSION

This section provides conclusions on the different aspects of this research by answering the research questions posed in the introduction of this paper.

The first research question posed is "*How can Footprint and Capabilities be defined?*". As described in section "*V. Downstream CDN selection criteria*", the general idea of the IETF CDNi Working Group about Footprint and Capabilities is that an uCDN can make an initial decision for a certain dCDN by looking at the Footprint information, whereas additional Capabilities information can be used when the Footprint information is insufficient to make a delegation decision. However after discussing and analysing the definition and purpose of both terms we can conclude that the clear distinction made between both terms is not valid.

A better approach would be not to divide the selection process into these two stages, but to make the footprint information part of the capabilities requirements. Then the selection of the dCDN would not only be based on the footprint information but will be based on a selection of capabilities. In this case the selection could be much more sophisticated and situations as described in section "*V. Downstream CDN selection criteria*" could more easily be avoided.

Based on the definition of the Footprint and Capabilities information, the different protocols for exchanging the information have been described and analysed to answer the second research question "*Which proposed method is more suitable for exchanging footprints and capabilities between different CDNs?"*.

As the extension on M-BGP for CDNi has some more drawbacks in comparison with the ALTO protocol, the ALTO protocol seems to be a better candidate for exchanging footprint as well as capabilities information. Although this protocol is based on a client-server architecture which could lead to single point of failure situations, the protocol also has some advantages compared to the extension on M-BGP for CDNi.

An advantage of the ALTO is that the protocol is at the time of writing still in the phase of definition and standardisation. This makes it more easy to suggest additional functionalities that can be used within the CDNi as the protocol is still not fully defined. Also in comparison with M-BGP, where nodes on the path are not obligated to understand all options, there is no downside of only implementing the ALTO protocol within different CDNs as the protocol does not rely on other entities that should understand the protocol except for the client and the server. Besides that, the framework of the ALTO protocol is defined in a more flexible way to support future extensions.

So of the different protocols suggested by the IETF CDNi Working Group and analysed in this paper, the ALTO protocol seems to be the better candidate for exchanging footprint information between CDNs at this time.

IX. FUTURE RESEARCH

As the CDNi framework is still under discussion by the IETF CDNi Working Group, there are many components still unclear as also not yet fully defined. This leaves enough room for future research into the different components of the framework as also the corresponding processes. An interesting question could be for example why the IETF CDNi Working Group has defined four different interfaces within the framework and not for example three different interfaces as currently defined within the ETSI standardisation process for interconnecting CDNs.

Another interesting point of research could be a decentralised version of the ALTO protocol to try to mediate the single point of failure cases the currently being defined ALTO protocol can be vulnerable of. Also other Application Layer protocols could be compared to see whether they are capable of exchanging the information we defined as Footprint and Capabilities information. An example of one of these protocol is the by the IETF CDNi Working Group recently suggested Software-defined networking [41].

Once the ALTO protocol has been better defined it is interesting to look into the possibilities to setup a sort of proof

of concept to see whether the conclusions and suggestion made in this paper as also by the IETF CDNi Working Group are correct.

REFERENCES

[1] A. Vakali, and G. Pallis, “Content Delivery Networks: Status and Trends,” IEEE Internet Computing, IEEE Computer Society, pp. 68-74, November-December 2003.

[2] B. Krishnamurthy, C. Willis, and Y. Zhang, “On the Use and Performance of Content Distribution Network,” In Proceedings of 1st International Internet Measurement Workshop, ACM Press, pp. 169-182, 2001.

[3] G. Pallis, and A. Vakali, “Insight and Perspectives for Content Delivery Networks,” Communications of the ACM, Vol. 49, No. 1, ACM Press, NY, USA, pp. 101-106, January 2006..

[4] A.M.K. Pathan, and R. Buyya, “A Taxonomy and Survey of Content Delivery Networks”, Grid Computing and Distributed Systems Laboratory, The University of Melbourne, Australia. February, 2007

[5] N. Bartolini, E. Casalicchio, and S. Tucci, “A Walk Through Content Delivery Networks,” In Proceedings of MASCOTS 2003, LNCS Vol. 2965/2004, pp. 1-25, April 2004.

[6] What is Global Server Load Balancing | How GSLB works | Load Balance Dedicated Servers, http://kb.eukhost.com/global-server-load-balancing/, June 2012, status [online]

[7] J. Dilley, B. Maggs, J. Parikh, H. Prokop, R. Sitaraman, and B. Weihl, “Globally Distributed Content Delivery,” IEEE Internet Computing, pp. 50-58, September/October 2002.

[8] CDNi Info Page, https://www.ietf.org/mailman/listinfo/cdni, June 2012, status [online]

[9] T. Bates, R. Chandra, D. Katz and Y. Rekhter, “Multiprotocol Extensions for BGP-4”, Internet Engineering Task Force, RFC4760, January 2007, http://tools.ietf.org/rfc/rfc4760.txt

[10] H. Ould-Brahim, D. Fedyk and Y. Rekhter, “BGP Traffic Engineering Attribute”, Internet Engineering Task Force, RFC 5543, May 2009, http://tools.ietf.org/rfc/rfc5543.txt

[11] P. Mohapatra, R. Fernando, E.C. Rosen and J. Uttaro, “The Accumulated IGP Metric Attribute for BGP”, Internet Engineering Task Force, draft-ietf-idr-aigp-08 , June 2012, http://tools.ietf.org/id/draft-ietf-idr-aigp-08.txt

[12] J. Seedorf and E. Burger, “Application-Layer Traffic Optimization (ALTO) Problem Statement”, Internet Engineering Task Force, RFC5693, October 2009, http://tools.ietf.org/rfc/rfc5693.txt

[13] R. Alimi, R. Penno and Y. Yang, “ALTO Protocol”, Internet Engineering Task Force, draft-ietf-alto-protocol-11, March 2012, http://tools.ietf.org/id/draft-ietf-alto-protocol-11.txt

[14] S. Kiesel, M. Stiemerling, N. Schwan, M. Scharf and H. Song, “ALTO Server Discovery”, Internet Engineering Task Force, draft-ietf-alto-server-discovery-03, March 2012, http://tools.ietf.org/id/draft-ietf-alto-server-discovery-03.txt

[15] S. Kiesel, S. Previdi, M. Stiemerling, R. Woundy and R. Yang, “Application Layer Traffic Optimization (ALTO) Requirements”, Internet Engineering Task Force, draft-, June 2012, http://tools.ietf.org/id/draft-ietf-alto-reqs-16.txt

[16] M. Stiemerling, S. Kiesel and S. Previdi, “ALTO Deployment Considerations”, Internet Engineering Task Force, draft-ietf-alto-deployments-04, March 2012, http://tools.ietf.org/id/draft-ietf-alto-deployments-04.txt

[17] B. Niven-Jenkins, G. Watson, N. Bitar, J. Medved and S. Previdi, “Use cases for ALTO within CDNs”, Internet Engineering Task Force, draft-jenkins-alto-cdn-use-cases-03, Juni 2012, http://tools.ietf.org/id/draft-jenkins-alto-cdn-use-cases-03.txt

[18] ALTO status page, http://tools.ietf.org/wg/alto/, June 2012, status [online]

[19] G. Bertrand, F. Le Faucheur and L. Peterson, “Content Distribution Network Interconnection (CDNI) Experiments”, Internet Engineering Task Force, draft-bertrand-cdni-experiments-02, February 2012, , http://tools.ietf.org/id/draft-bertrand-cdni-experiments-02.txt

[20] Stef van der Ziel, http://www.jet-stream.com/stefvanderziel/

[21] B. Molina, C.E. Palau, M. Esteve and J. Lloret, “On Content Delivery Network protocols and applications”, Communication Department, Polytechnical University of Valencia, Spain, 2004

[22] AS Names, http://bgp.potaroo.net/cidr/autnums.html, June 2012, status [online]

[23] Y. Rekhter, T. Li and S. Hares, “A Border Gateway Protocol 4 (BGP-4)”, Internet Engineering Task Force, RFC4271, January 2006, http://tools.ietf.org/rfc/rfc4271.txt

[24] S. Sangli, D. Tappan and Y. Rekhter, “BGP Extended Communities Attribute”, Internet Engineering Task Force, RFC4360, February 2006, http://tools.ietf.org/rfc/rfc4360.txt

[25] T. Bates, R. Chandra, D. Katz and Y. Rekhter, “Multiprotocol Extensions for BGP-4”, Internet Engineering Task Force, RFC 4760, January 2007, http://tools.ietf.org/rfc/rfc4760.txt

[26] B. Niven-Jenkins, F. Le Faucheur, and N. Bitar, “Content Distribution Network Interconnection (CDNI) Problem Statement”, Internet Engineering Task Force, draft-ietf-cdni-problem-statement-06, May 2012, http://tools.ietf.org/id/draft-ietf-cdni-problem-statement-06.txt

[27] G. Bertrand, E. Stephan, T. Burbridge, P. Eardley, K. Ma, and G. Watson, “Use Cases for Content Delivery Network Interconnection”, Internet Engineering Task Force, draft-ietf-cdni-use-cases-08, June 2012. http://tools.ietf.org/id/draft-ietf-cdni-use-cases-08.txt

[28] P. Rzewski, M. Day, and D. Gilletti, “Content Internetworking (CDI) Scenarios”, Internet Engineering Task Force, RFC 3570, July 2003. http://tools.ietf.org/pdf/rfc3570.pdf

[29] L. Peterson, and B. Davie, “Framework for CDN Interconnection, Internet Engineering Task Force, draft-ietf-cdni-framework-00, April 2012. http://tools.ietf.org/id/draft-ietf-cdni-framework-00.txt

[30] J. Seedorf, J. Peterson and S. Pervidi, “CDNI Request Routing: Footprint and Capabilities Semantics”, Internet Engineering Task Force, draft-spp-cdni-rr-foot-cap-semantics-00, March 2012. http://tools.ietf.org/id/draft-spp-cdni-rr-foot-cap-semantics-00.txt

[31] X. He, S. Dawkins, G. Chen, Y. Zhang and W. Ni, “Capability Information Advertising for CDN Interconnection”, Internet Engineering Task Force, draft-he-cdni-cap-info-advertising-01, March 2012. http://tools.ietf.org/id/draft-he-cdni-cap-info-advertising-01.txt

[32] K. Leung and Y. Lee, “Content Distribution Network Interconnection (CDNI) Requirements”, Internet Engineering Task Force, draft-ietf-cdni-requirements-03, June 2012. http://tools.ietf.org/id/draft-ietf-cdni-requirements-03.txt

[33] J. Hawkinson and T. Bates, “Guidelines for creation, selection, and registration of an Autonomous System (AS)”, Internet Engineering Task Force, RFC1930, March 1996. http://tools.ietf.org/rfc/rfc1930.txt

[34] "ISO - Maintenance Agency for ISO 3166 country codes - What is ISO 3166?", International Organization for Standardization, http://www.iso.org/iso/country_codes/background_on_iso_3166/what_is_iso_3166.htm

[35] G. Bertrand, "CDN Footprint Discovery", Internet Engineering Task Force, draft-bertrand-cdni-footprint-discovery-00, March 2012, http://tools.ietf.org/id/draft-bertrand-cdni-footprint-discovery-00.txt

[36] X. He, S. Dawkins, G. Chen, W. Ni and Y. Zhang, "Routing Request Redirection for CDN Interconnection", Internet Engineering Task Force, draft-he-cdni-routing-request-redirection-01.txt, February 2012, http://tools.ietf.org/id/draft-he-cdni-routing-request-redirection-01.txt

[37] K. Ma, "Content Distribution Network Interconnection (CDNI) Metadata Interface", Internet Engineering Task Force, draft-ma-cdni-metadata-02, April 2012, http://tools.ietf.org/id/draft-ma-cdni-metadata-02.txt

[38] S. Pervidi, F. Le Faucheur, J. Medved and A. Guillou, "CDNI Footprint Advertisement", Internet Engineering Task Force, draft-previdi-cdni-footprint-advertisement-01, March 2012, http://tools.ietf.org/id/draft-previdi-cdni-footprint-advertisement-01.txt

[39] J. Seedorf, "CDNI Request Routing with ALTO, Internet Engineering Task Force, draft-seedorf-cdni-request-routing-alto-01, March 2012, http://tools.ietf.org/id/draft-seedorf-cdni-request-routing-alto-01.txt

[40] E. Stephan and S. Ellouze, "ALTO extensions for CDNi", Internet Engineering Task Force, draft-stephan-cdni-alto-session-ext-00, March 2012, http://tools.ietf.org/id/draft-stephan-cdni-alto-session-ext-00.txt

[41] M-K. Kim, H-J. Kim, D. Chang and T.Kwon, "CDNI Request Routing with SDN", Internet Engineering Task Force, draft-shin-cdni-request-routing-sdn-00, July 2012, http://tools.ietf.org/id/draft-shin-cdni-request-routing-sdn-00.txt

[42] Stef van der Ziel, personal communication, June 15, 2012

[43] J. Scudder and R. Chandra, "Capabilities Advertisement with BGP-4", Internet Engineering Task Force, RFC5492, February 2009, http://tools.ietf.org/rfc/rfc5492.txt