CDNI G. Chen
Internet-Draft China Telecom
Intended status: Informational M. Li
Expires: December 30, 2012 H. Xia
ZTE Corporation
J. Liang
China Telecom
June 28, 2012

Intra-CDN Provider CDNi Experiment
draft-chen-cdni-intra-cdn-provider-cdni-experiment-00

Abstract

In [I-D.ietf-cdni-use-cases], the Inter-Affiliates CDN Interconnection use case is described. In this scenario, a large CDN Provider may have several autonomous or semi-autonomous subsidiaries that each operates on their own CDN. The CDN Provider needs to make these down-stream CDNs interoperate to provide a consistent service to its customers on the whole collective footprint.

This document illustrates in details the CDNi experiment that has been carried out by China Telecom, and the lessons and experiences to CDNi standardization work.

Table of Contents

1. Introduction

As a CDN service provider, China Telecom has established video CDNs in more than ten provinces in China. These video CDNs, provided by different vendors, are relatively independent and only provide services to the end users of their own provinces. Under this circumstance, if a Content Provider (CP) wants to provide services to multiple provinces, it needs to interact with CDNs in other provinces via interfaces which may support different standards.

China Telecom launched the CDN interconnection trial network in 2011 where CDNs from six different vendors (ZTE, Huawei, Cisco, etc.) were used to conduct the interconnection experiment in three provinces. This experiment aims at testing the scenario where the operator provides autonomous services via CDN interconnection in order to provide enhanced user experience. It is noted that a simplification of the interconnection framework and the corresponding procedures would really improve the service and real-time viewing experience.

This experiment is not intended to cover all of the use cases or the scenarios that are within the scope of CDNi work. It simply provides some practical information gathered from the actual network experiment as a reference for the CDNi standardization work. These CDN interconnection implementation experiments cover mostly the intra-operator interconnection scenarios.

1.1. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

This document reuses the terminology defined in:

[I-D.draft-ietf-cdni-problem-statement-06],

[I-D.draft-ietf-cdni-requirements-03],

[I-D.draft-ietf-cdni-framework-00], and

[I-D.draft-ietf-cdni-use-cases-08].

2. Intra-CDN Provider CDNi Experiments

2.1. Experiment Configuration

The interconnection of four CDNs in two provinces has been tested in this experiment. Each province has two CDNs interconnected which are provided by different CDN vendors. As depicted in Figure 1, CDN A1 of Province A has contracts with service providers CP1 and CP2, and it acts as the content storage center of the nation.CDN B1 of Province B has contract with service provider CP3. Meanwhile, CDN A1 and CDN B1 are the sub-center CDNs of respective province, while CDN A2 and CDN B2 are the regional CDNs of respective province. CDN A1 and CDN B1 are deployed on the provincial backbone networks, while CDN A2 and CDN B2 are deployed on the MANs. CDN A1 is the upstream CDN of CDN A2. Services are provided to the end users by certain node in regional CDN, which is usually the geographically closest one to the end user. However, if this desired node is overloaded, certain re-routing or load-balancing criteria could be used to choose another node. CDN B1 and CDN B2 have similar deployment. The provincial center CDN A1 and CDN B1 are interconnected with each other. They do not have interconnection with the region CDNs in other provinces than themselves. The regional CDNs of the respective provinces do not interconnect with one another either.

China Telecom's CDN trial network offers two types of services: intra-province service (provided by CP2 and CP3) and inter-province service (provided by CP1). Intra-province service is provided independently within the province without any interconnection with CDNs in other provinces. When inter-province service is provided, content is ingested to the CDN in a single province and then distributed among the CDNs in all other provinces in the trial network. In this experiment the inter-province service is ingested via CDN A1 node and the end users can obtain services through the CDNs that are located in their own provinces.

```
                    +------+
                    |  CP1 |
                    +------+             +-----+
     +------+          /                 | CP3 |
     |  CP2 |         /                  +-----+
     +------+        /        :              /
         \          /         :             /
          \        /          :            /
           \                  :
              _,.---.,,       :         _,.---.,,
            .`         `.     :       .`         `.
           '             \    :      '             \
          |    CDN A1     |---:---- |     CDN B1    |
           ,             / ---:----  ,             /
            ',         ,-     :       ',         ,-
              ``''--'``       :         ``''--'``
                | |           :            | |
                | |           :            | |
                | |           :            | |
              _,.---.,,       :         _,.---.,,
            .`         `.     :       .`         `.
           '             \    :      '             \
          |    CDN A2     |   :     |     CDN B2    |
           ,             /    :      ,             /
            ',         ,-     :       ',         ,-
              ``''--'``       :         ``''--'``
                              :
                              :
                  Pronvince A :  Province B
                              :
```

Figure 1 CDNI between Two Different Provinces, each with Two CDNs

The details of the experiment are as presented below:

CP3 has contract with CDN B1 to provide content delivery service, e.g., IPTV service, to the end users in the Province B region. CP3 does not serve end users outside Province B. As an autonomous service by China Telecom, the content of this service is ingested into CDN B1. As its downstream CDN, CDN B1 can delegate the content requests from end users to CDN B2 to perform content delivery.

When CDN B1 receives the content request from EU B of Province B related to service provided by CP3, it redirects this request to CDN B2. If CDN B2 has locally cached the copy of the content requested by EU B (cache hit case), it serves EU B directly by delivering the content to EU B. If CDN B2 does not have the content cached (cache miss case), it acquires the content from CDN B1.

CP1 has contract with CDN A1 to provide content delivery service, e.g., OTT service, to end users within Province A and in the region of Province B. The content for this service is ingested from CDN A1. As for the cross-province routing, since it is within the same operator, static configuration can be used for dCDN selection. In this experiment, the national content storage center CDN A1 configures locally the relationship table of the end user's IP addresses and the loading condition of dCDNs.

When CDN A1 receives a content request from EU B of Province B, it redirects the request to CDN B2 directly after it checks the local configuration table without going through multiple redirection processes like CDN A1->CDN B1->CDN B2. If CDN B2 has locally cached a copy of the content that has been requested by EU B (cache hit case), it serves EU B directly by delivering the content to EU B. If CDN B2 does not have the content cached (cache miss case), it acquires the content from CDN B1. If CDN B1 does not have the content cached either, it acquires the content from CDN A1.

In this experiment, we use a content acquisition method that is different from the current CDNi work. The method is based on Content Identification by using UniContentID that is defined to uniquely identify a content item. A detailed description of this method is presented in Section 2.3.

```
                          +------+
                          |  CP1 |
                          +------+             +-----+
                             /                 | CP3 |
                            /                  +-----+
                           /        :              /
                          /         :             /
                         /          :            /
                                    :
                    _,.---.,,       :         _,.---.,,
                  .`         `.     :       .`         `.
                 '             \    :      '             \
                |    CDN A1     |---:---- |     CDN B1    |
                 ,             / ---:----  ,             /
                  ',         ,-     :       ',         ,-
                    ``''--'``       :         ``''--'``
                      | |           :            | |
                      | |           :            | |
                      | |           :            | |
                    _,.---.,,       :         _,.---.,,
                  .`         `.     :       .`         `.
                 '             \    :      '             \
                |    CDN A2     |   :     |     CDN B2    |
                 ,             /    :      ,             /
                  ',         ,-     :       ',         ,-
                    ``''--'``       :         ``''--'``
                                    :             |
                         Province A :  Province B |
                                    :          +------+
                                    :          | EU B |
                                    :          +------+
                                    :
                                    :
```

Figure 2 CDNI between Two Different Provinces

2.2. Logging

Since in this experiment CDN interconnection is implemented within the scope of the same operator, charging-related operations via Logging Interface are not required. Therefore, we have neither implemented nor tested any Logging operations.

2.3. UniContentID

In the current IETF CDNi standards, it is required to add the URL of original request and the URL in the process through CDNs in the request routing. When cache is not hit, downstream CDN needs to use the information above to trace to the source. The redirection flows

are complex and the URL becomes very long which makes the implementation very difficult.

In this experiment, the UniContentID as defined in [I-D.draft-chen-cdni-rr-content-acquisition] is used. UniConentID is described by two tuple as (ProviderID, ContentID), e.g. ('iptv.netitv.com','01234567890123456789012345678900'), which can uniquely identify a content item. We trace the content source according to the configuration table of ProviderID and the corresponding relationship between ProviderID and the IP address of upstream CDN. It is our view that redirection and content acquisition are different routes.

2.4. Request Routing and Content Acquisition

2.4.1. Request Routing and Content Acquisition in Province B

```
               +--------------------------+   +-------------------------+
               |        CDN B1            |   |        CDN B2           |
   +-------+   | +---------+ +----------+|   | +----------+  +---------+|
   |  EU   |   | |CDN B1 RR| | CDN B1 DN||   | |CDN B2 RR|  |CDN B2DN||
   +-------+   | +---------+ +----------+|   | +----------+  +---------+|
       |       +--------------------------+   +-------------------------+
       |(1)RTSP or HTTP REQ          |                 |            |
       |------------> |              |                 |            |
       | (2)RTSP or HTTP RES         |                 |            |
       | <------------|              |                 |            |
       |              | (3)RTSP or HTTP REQ            |            |
       |--------------------------------------------> |            |
       |              | (4)RTSP or HTTP REP            |            |
       | <---------------------------------------------|            |
       |              |              |                 |            |
       |              |  (5)RTSP or HTTP REQ           |            |
       |---------------------------------------------------------> |
       |              |              |                 |            |
       |              |              |      (6)HTTP REQ|            |
       |              |              |<--------------------------- |
       |              |              |                 |            |
       |              |              |      (7)HTTP REP|            |
       |              |              |---------------------------> |
       |              |              |                 |            |
       |              |              |                 |            |
       |              |              |                 |            |
       |              |        (8)DATA                 |            |
       |<--------------------------------------------------------- |
       |              |              |                 |            |
```

Figure 3 Request Routing and Content A Acquisition in Province B

The Message sequence of Figure 3 is shown below in details.

(1) End-User sends a request to the load balancer of CDN B1, i.e. RR of CDN B1 for the content.The URL includes the parameter of CMSID and Domain(Please refer to [I-D. draft-chen-cdni-rr-content-acquisition] for corresponding definitions).

(2) The RR of CDN B1 chooses an optimal RR of dCDN, i.e. RR of CDN B2 for the End-User according to the load of dCDN and the IP Pool information.

(3) End-User sends a request to the dCDN, i.e. RR of CDN B2 for content acquisition.

(4) RR of CDN B response to the End-User for the information of a delivery node i.e. DN.

(5) End-User sends a content request to the DN of CDN B2, the URL include the information of CMSID and Domain. The DN of CDN B2 analyses the ProviderID according the information of CMSID and Domain and looks up if the content exists in the cache according to the ProviderID and ContentID. If the content is cached, the DN of CDN B2 serves the End-User. Otherwise, it skips to step (6).

(6) The DN of CDN B2 looks up the configuration table and determines, according to the ProviderID, to acquire content uniquely identified by the ProviderID and ContentID from the DN of CDN B1.

(7) The content in the DN of CDN B1 relays to the DN of CDN B2.

(8) The relayed content is served by the DN of CDN B2 to the End-User via playing by downloading.

2.4.2. Request Routing and Content Acquisition between Province A and Province B

```
         +------------------+ +------------------+ +------------------+
         |        CDN A1    | |        CDN B1    | |        CDN B2    |
   +----+|+------+ +------+ | |+------+ +------+ | |+------+ +------+ |
   | EU |||CDN A1| |CDN A1| | ||CDN B1| |CDN B1| | ||CDN B2| |CDN B2| |
   +----+||  RR  | | DN   | | ||  RR  | | DN   | | ||  RR  | | DN   | |
     |   |+------+ +------+|| |+------+ +------+|| |+------+ +------+||
     |   +------------------+ +------------------+ +------------------+
     |(1)RTSP or HTTP REQ         |         |          |          |
     |-----> |         |          |         |          |          |
     |(2)RTSP or HTTP REP         |         |          |          |
     | <-----|         |          |         |          |          |
     |       |  (3)RTSP or HTTP REQ         |          |          |
     |----------------------------------------------->  |          |
     |       |  (4)RTSP or HTTP REP         |          |          |
     | <-----------------------------------------------|          |
     |       |         |          |         |          |          |
     |       |         |   (5)RTSP or HTTP REQ         |          |
     |----------------------------------------------------------->|
     |       |         |          |         |          |          |
     |       |         |          |         |   (6)HTTP REQ       |
     |       |         |          |         |<--------------------|
     |       |         |  (7)HTTP REQ       |          |          |
     |       |         |<-------------------|          |          |
     |       |         |          |         |          |          |
     |       |         |  (8)HTTP REP       |          |          |
     |       |         |------------------->|          |          |
     |       |         |          |         |    (9)HTTP REP      |
     |       |         |          |         |-------------------->|
     |       |         |          |         |          |          |
     |       |         |  (10)DATA|         |          |          |
     |<---------------------------------------------------------- |
     |       |         |          |         |          |          |
     |       |         |          |         |          |          |
```

Figure 4 RR and Content Acquisition between Province A and Province B

The Message sequence of Figure 4 is shown below in details.

Step (1)~(5) is similar to step(1)~(5)of Section 2.4.1.Note that the request routing process does not need the participation of CDN B1.

(6) The DN of CDN B2 looks up the configuration table and determines, according to the ProviderID, to acquire content uniquely identified by the ProviderID and ContentID from the DN of CDN B1. If the content is cached in DN of CDN B1, the DN of CDN B1 serves the End-User. Otherwise, it skips to step (7).

(7) The DN of CDN B1 looks up the configuration table and determines, according to the ProviderID, to acquire content from the DN of CDN

A1.

(8) The content in the DN of CDN A1 relays to the DN of CDN B1.

(9) The content in the DN of CDN B1 relays to the DN of CDN B2.

(10) The relayed content is served by the DN of CDN B2 to the End-User via playing for downloading.

2.4.3. Test Results

Based on the experiment model above, we have tested the CDN interconnection scenario in China Telecom's trial network where 1 Gbps video traffic is not hit in the local cache and content acquisition is needed. Performance tests are done by using the tools from Shineck and Spirent. We have made such operations as fast forward, fast rewind and positioning play etc. The testing results show that the response time is less than one second and the Max DF is less than 50msec. Also we conducted the test for a period of six hours for stability tests through complex operation of fast forward, fast rewind, positioning play, etc. The tests results show that the response time is less than one second and call loss is less that 0.1%. During the process of performance tests, the user requests the video on demand and Live contents through set-up box and conducts complex operation such as fast forward, fast rewind, positioning play, etc. The program is smooth during the play. The tests show that the CDN interconnection architecture for intra-operator video service is both feasible and efficient.

2.5. Control

In the IETF CDNi draft [I-D.murray-cdni-triggers], the uCDN controls the content operation, i.e., content adding, content purging and content modifying are performed through the control interface. In this section, we describe a general content control flow by using CDN B1 and CDN B2 as an example which are used in this experiment.

```
+--------+                +--------+
| CDN B1 |                | CDN B2 |
+--------+                +--------+
    |                          |
    |HTTP://dCDN IP:PORT/ContentDeployReq
    |-----------------------> |
    |                          |
    |ContentDeployReqResponse |
    |<----------------------- |
    |                          |
    |                          |
    |Achieve XML through FTP  |
    |<----------------------- |
    |                          |
    |                          |
    |                          |
    |http://uCDN IP:PORT/ContentDeployResult
    |<----------------------- |
    |                          |
    |                          |
    |                          |
    |ContentDeployResultResponse
    |-----------------------> |
    |                          |
    |                          |
```

Figure 5 Message Exchange for Content Acquisition

The Message sequence of Figure 5 is shown below in details.

(1) CDN B1 sends a content management requests to the CDN B2 including the content adding, content purging and content modifying. The content object can be live content, video on demand content or TV on demand. The object of ContentDeployReq includes the URL address of XML description of the content object.

(2) CDN B2 checks if the URL FTP address from CDN B1 is OK. If the result is OK, CDN B2 responds positively (success) to the CDN B1.

(3) CDN B2 login in the FTP of CDN B1 to achieve the XML data of content object and executes the operation according to the instruction of CDN B1, i.e., content adding, content purging or content modifying.

(4) CDN B2 responds to the CDN B1 for the operation results, i.e., with either a success or a failure result.

(5) CDN B1 confirms to the CDN B2 for the operation result and registers the operation result.

3. Lessons Learned

The CDNi functionality tested in this experiment is applicable to intra-operator case only. The inter-province service is ingested via CDN in one province and can be used throughout the entire trial network in two provinces. During the initial stage of CDNi standardization, the most practical scenario to be considered is the interconnection of CDNs distributed in different geographical regions within one operator so as to enhance the consistency and continuity of the services provided by operators themselves. Therefore, this experiment is intended to provide some experiences and references on CDNi networking to those operators who have initial requirements for internal CDN interconnection across different geographical regions.

3.1. Simplification of operation procedures

It is demonstrated that relatively simple methods can be used to simplify or optimize the Request Redirection, Content Acquisition, Content Pre-Positioning, Content Addition/Modification/Deletion procedures. This is conducive to operators when they also act as service providers to serve the end users by fulfilling the requirements of of real-time service and demanding user experience.

3.2. Redirection

According to the HTTP/DNS-based Request Routing process defined in the current [I-D.ietf-cdni-framework], multiple redirection processes are needed to determine the final CDN node that is suitable to serve the end user. This may be convenient in the case of two-level CDNs. But in case of large-scale CDN networking or complex CDN topology, this would cause serious delay. The method provided in this experiment, i.e., the one based on the local configuration of the relationship table of end users' IP addresses and load situation of dCDNs, the uCDN, as the national content storage center, can quickly acquire and locate the final dCDN that is suitable to serve the end user. By usingthis technique, the routing selection can be largely simplified especially when operators have large-scale internal networking.

3.3. UniContentID

In current IETF CDNi work [I-D.ietf-cdni-framework], the content acquisition by dCDN from uCDN is achieved via embedding the URL of the original request as well as that of the CDN the request is redirected to during the redirection process. This would result in a very long URL. Limited by the length and format of URL, such approach would cause serious delay and waste of resources. In addition, it would be difficult to implement this, especially under

the complex CDNi topology.

The unique content identification (named UniContentID in this document) can be used to uniquely identify the content item ingested by the content provider. The content source can also be resolved from UniContentID contained in the end user's content request for content acquisition. Due to the uniformity of UniContentID format and its unchangeable nature during the transmission among CDNs, it can all along identify the content requested by the end user even after multiple forwards under complex CDNi topology. Note that these introduce the need for defining a unique content identification for content item in CDNi framework.

3.4. Metadata and Logging

The metadata related to CDNi content delivery would be relatively simple. Such policy information as content ingestion, content acquisition, etc. can be achieved by pre-configuration. It is also due to the scope within one single operator, there is no need for charging-related operations via logging interface, which, as a result, can be simplified or may not even be supported.

3.5. Inter-Operator CDN Interconnection

For CDN interconnection across operators, if the operators are clearly aware of each other's CDN framework (according to their agreements), the method that is used in this experiment can also be utilized as a reference, i.e., for implementing end user's request redirection and content acquisition via maintaining the configuration table, which can also achieve relatively high content delivery efficiency. For those operators who have complicated internal CDN topology or proprietary APIs or CDN topology, it is our view that it requires to seek solutions for dynamic request redirection and content acquisition.

4. Security Considerations

This experiment is carried out within a single operator. No security issues are considered at this stage of the experiment.

5. IANA Considerations

This memo has no IANA Considerations.

6. Acknowledgments

To be added later

7. References

7.1. Normative References

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

7.2. Informative References

[I-D.ietf-cdni-framework]
Peterson, L. and B. Davie, "Framework for CDN Interconnection", April 2012.

[I-D.ietf-cdni-problem-statement]
Niven-Jenkins, B., Faucheur, F., and N. Bitar, "Content Distribution Network Interconnection (CDNI) Problem Statement", May 2012.

[I-D.ietf-cdni-requirements]
Leung, K. and Y. Lee, "Content Distribution Network Interconnection (CDNI) Requirements", December 2011.

[I-D.ietf-cdni-use-cases]
Bertrand, G., Stephan, E., Burbridge, T., Eardley, P., Ma, K., and G. Watson, "Use Cases for Content Delivery Network Interconnection", June 2012.

[I-D.murray-cdni-triggers]
Murray, R. and B. Niven-Jenkins, "CDN Interconnect Triggers", February 2012.

[I-D.narten-iana-considerations-rfc2434bis]
Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", draft-narten-iana-considerations-rfc2434bis-09 (work in progress), March 2008.

[RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", RFC 2629, June 1999.

[RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, July 2003.

Authors' Addresses

Ge Chen
China Telecom
109 West Zhongshan Ave
Guangzhou, Tianhe District
China

Phone:
Email: cheng@gsta.com

Mian Li
ZTE Corporation
Nanjing,   210012
China

Phone:
Email: li.mian@zte.com.cn

Hongfei Xia
ZTE Corporation
Nanjing,   210012
China

Phone:
Email: xia.hongfei@zte.com.cn

Jie Liang
China Telecom
109 West Zhongshan Ave
Guangzhou, Tianhe District
China

Phone:
Email: liangj@gsta.com