CDNi Request Routing Redirection with Loop Prevention

draft-choi-cdni-req-routing-redir-loop-prevention-01.txt

Taesang Choi (choits@etri.re.kr)
Young-IL Seo (yohan.seo@kt.com)
Dong-Ju Kim (dj.kim@kt.com)
Jongmin Lee (jminlee@sk.com)
Ja-Ryeong Koo (wjbkoo@lguplus.co.kr)
John Dongho Shin (eastsky@solbox.com)
Kyungsoo Park (kspark@kaist.ac.kr)

November 8, 2012
85th IETF, Atlanta
CDNi WG

Update Overview

- Presented at 84th IETF, Vancouver
- First revision based on comments made at 84th IETF
  - Comments on Content-Provider-ID formats
  - Comment on URL length limitation
  - Comment on loop prevention algorithm
- Experimentations & results

Main Changes (1)

- In 00 version, "CDN-Provider-ID“ was described as a list of CDN provider Names and MaxNumRedHops
- In 01 version, changed to a list of CDN-Provider-Names followed by MaxNumRedHops.
- Note that a list of CDN-Provider-ID is conveyed in URI string to deal with HTTP URL length limitation
- Example: http://dcdn1.csp.com/ucdn.csp.com?uCDN=100:0 & dDCN=200:1 & dCDN=300:0 & MaxNumRedHops=8

Main Changes (2)

- In 00 version, we specified loop prevention algorithm in pseudo code
- In 01 version, we changed it to specify the following in descriptive form:
  - a mechanism to allow loop detection
  - post processing, that is, who is responsible and in what quality (service availability vs quality) for resolving the situation

Experimentation of Loop Prevention

- Built a PoC test-bed with our consortium members (KT, SKT, LGU+, SolBox)
  - Tested in a simple ring type cascaded topology
- Implemented both Iterative HTTP-/DNS-based request routing redirection. Recursive method is under way
- Objectives
  - Verify the feasibility of the proposed method
  - Measure delays incurred during RRR
  - Impact on the size and transmission performance of redirection messages

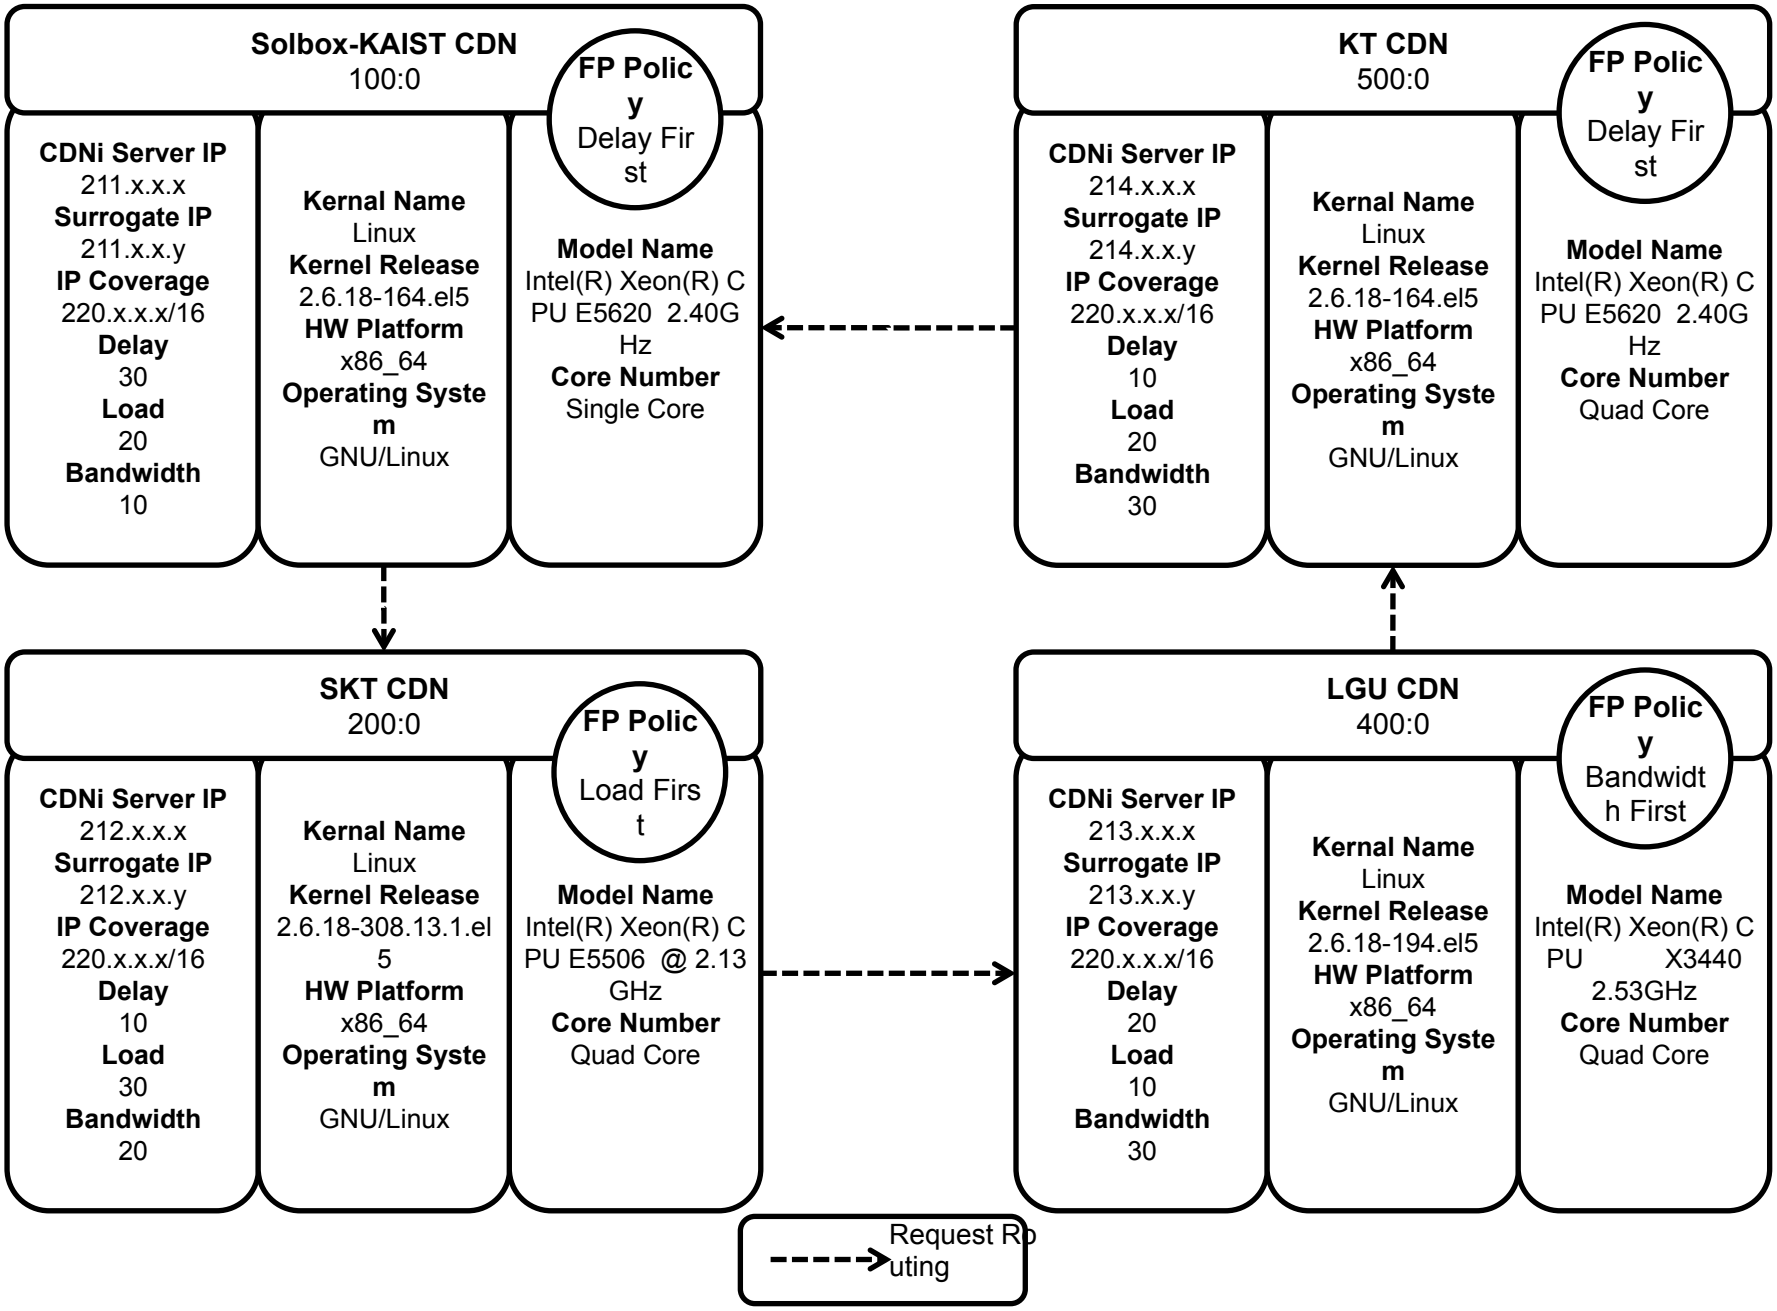
Solbox-KAIST CDN
100:0
FP Policy
Delay First
CDNi Server IP
211.x.x.x
Surrogate IP
211.x.x.y
IP Coverage
220.x.x.x/16
Delay
30
Load
20
Bandwidth
10
Kernal Name
Linux
Kernel Release
2.6.18-164.el5
HW Platform
x86_64
Operating System
GNU/Linux
Model Name
Intel(R) Xeon(R) CPU E5620 2.40GHz
Core Number
Single Core
KT CDN
500:0
FP Policy
Delay First
CDNi Server IP
214.x.x.x
Surrogate IP
214.x.x.y
IP Coverage
220.x.x.x/16
Delay
10
Load
20
Bandwidth
30
Kernal Name
Linux
Kernel Release
2.6.18-164.el5
HW Platform
x86_64
Operating System
GNU/Linux
Model Name
Intel(R) Xeon(R) CPU E5620 2.40GHz
Core Number
Quad Core
SKT CDN
200:0
FP Policy
Load First
CDNi Server IP
212.x.x.x
Surrogate IP
212.x.x.y
IP Coverage
220.x.x.x/16
Delay
10
Load
30
Bandwidth
20
Kernal Name
Linux
Kernel Release
2.6.18-308.13.1.el5
HW Platform
x86_64
Operating System
GNU/Linux
Model Name
Intel(R) Xeon(R) CPU E5506 @ 2.13GHz
Core Number
Quad Core
LGU CDN
400:0
FP Policy
Bandwidth First
CDNi Server IP
213.x.x.x
Surrogate IP
213.x.x.y
IP Coverage
220.x.x.x/16
Delay
20
Load
10
Bandwidth
30
Kernal Name
Linux
Kernel Release
2.6.18-194.el5
HW Platform
x86_64
Operating System
GNU/Linux
Model Name
Intel(R) Xeon(R) CPU X3440 2.53GHz
Core Number
Quad Core
Request Routing

Experimentation Findings

- Delays in various hop count settings: 5, 10, …, n hops
  - Iterative vs recursive
  - HTTP vs DNS
- Impact on the size and transmission performance
- Miscellaneous
  - 302 HTTP redirection supports upto 20 redirections
  - DNS CNAME supports upto 38 redirections
  - DNS redirection which retains initiating CDN (uCDN) domain name doesn't work. Work-around: replace with the immediate parent domain name instead

Delay graph

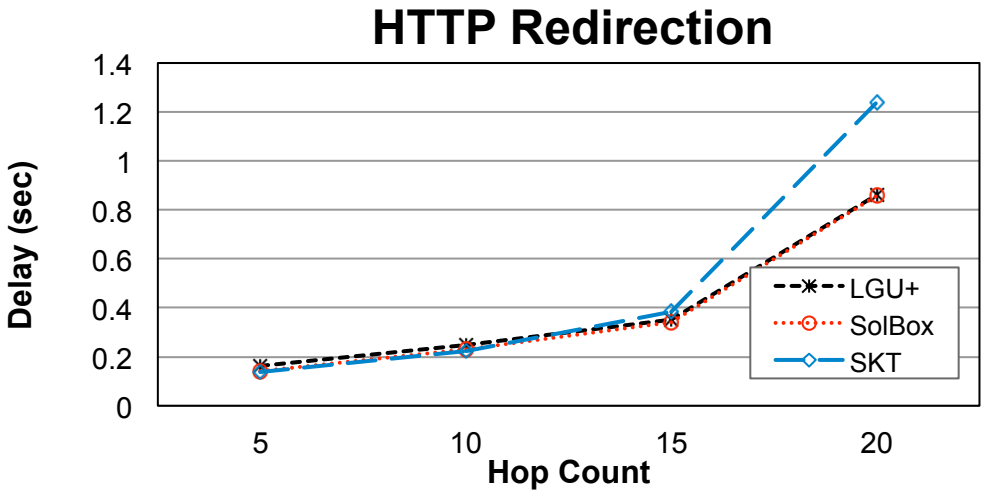
HTTP Redirection
Delay (sec)
1.4
1.2
1
0.8
0.6
0.4
0.2
0
5
10
15
20
Hop Count
LGU+
SolBox
SKT

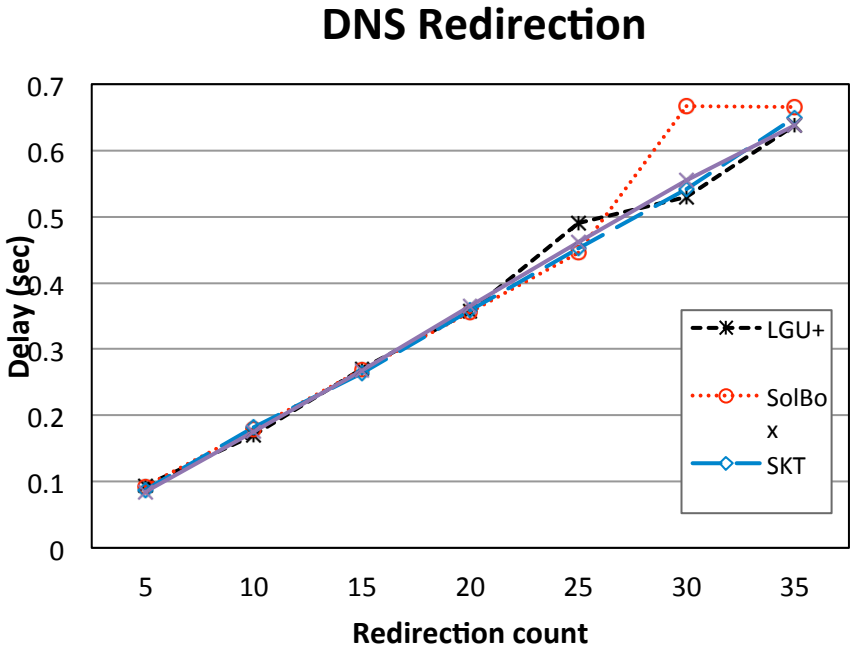
DNS Redirection
Delay (sec)
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
5
10
15
20
25
30
35
Redirection count
LGU+
SolBox
SKT

HTTP Redirections Trace

```
yonghwanbang@Kyle-mac:~$ wget http://solbox-kaistcdni.net
--2012-11-03 19:57:54--  http://solbox-kaistcdni.net/
Resolving solbox-kaistcdni.net (solbox-kaistcdni.net)... 211.38.137.37
Connecting to solbox-kaistcdni.net (solbox-kaistcdni.net)|211.38.137.37|:80... connected.
HTTP request sent, awaiting response... 302 Found
Location: http://sbcdn.skt-kaistcdni.net/solbox-kaistcdni.net/?ucdn=100:0&hopcnt=9 [following]
--2012-11-03 19:58:04--  http://sbcdn.skt-kaistcdni.net/solbox-kaistcdni.net/?ucdn=100:0&hopcnt=9
Resolving sbcdn.skt-kaistcdni.net (sbcdn.skt-kaistcdni.net)... 1.237.57.3
Connecting to sbcdn.skt-kaistcdni.net (sbcdn.skt-kaistcdni.net)|1.237.57.3|:80... connected.
HTTP request sent, awaiting response... 302 Found
Location: http://sktcdn.lgu-kaistcdni.net/solbox-kaistcdni.net/?ucdn=100:0&dcdn=200:0&hopcnt=8 [following]
--2012-11-03 19:58:14--  http://sktcdn.lgu-kaistcdni.net/solbox-kaistcdni.net/?ucdn=100:0&dcdn=200:0&hopcnt=8
Resolving sktcdn.lgu-kaistcdni.net (sktcdn.lgu-kaistcdni.net)... 211.115.90.78
Connecting to sktcdn.lgu-kaistcdni.net (sktcdn.lgu-kaistcdni.net)|211.115.90.78|:80... connected.
HTTP request sent, awaiting response... 302 Found
Location: http://lgucdn.solbox-kaistcdni.net/solbox-kaistcdni.net/?ucdn=100:0&dcdn=200:0&dcdn=400:0&hopcnt=7 [following]
--2012-11-03 19:58:24--  http://lgucdn.solbox-kaistcdni.net/solbox-kaistcdni.net/?ucdn=100:0&dcdn=200:0&dcdn=400:0&hopcnt=7
Resolving lgucdn.solbox-kaistcdni.net (lgucdn.solbox-kaistcdni.net)... 211.38.137.37
Connecting to lgucdn.solbox-kaistcdni.net (lgucdn.solbox-kaistcdni.net)|211.38.137.37|:80... connected.
HTTP request sent, awaiting response... 302 Found
Location: http://sbcdn.skt-kaistcdni.net/solbox-kaistcdni.net/?ucdn=100:0&dcdn=200:0&dcdn=400:0&dcdn=100:0&hopcnt=6 [following]
--2012-11-03 19:58:38--  http://sbcdn.skt-kaistcdni.net/solbox-kaistcdni.net/?ucdn=100:0&dcdn=200:0&dcdn=400:0&dcdn=100:0&hopcnt=6
Connecting to sbcdn.skt-kaistcdni.net (sbcdn.skt-kaistcdni.net)|1.237.57.3|:80... connected.
HTTP request sent, awaiting response... 302 Found
Location: http://sktcdn.lgu-kaistcdni.net/solbox-kaistcdni.net/?ucdn=100:0&dcdn=200:0&dcdn=400:0&dcdn=100:0&dcdn=200:0&hopcnt=5 [following]
--2012-11-03 19:58:38--  http://sktcdn.lgu-kaistcdni.net/solbox-kaistcdni.net/?ucdn=100:0&dcdn=200:0&dcdn=400:0&dcdn=100:0&dcdn=200:0&hopcnt=5
Connecting to sktcdn.lgu-kaistcdni.net (sktcdn.lgu-kaistcdni.net)|211.115.90.78|:80... connected.
HTTP request sent, awaiting response... 302 Found
Location: http://lgucdn.solbox-kaistcdni.net/solbox-kaistcdni.net/?ucdn=100:0&dcdn=200:0&dcdn=400:0&dcdn=100:0&dcdn=200:0&dcdn=400:0&hopcnt=4 [following]
--2012-11-03 19:58:38--  http://lgucdn.solbox-kaistcdni.net/solbox-kaistcdni.net/?ucdn=100:0&dcdn=200:0&dcdn=400:0&dcdn=100:0&dcdn=200:0&dcdn=400:0&hopcnt=4
Connecting to lgucdn.solbox-kaistcdni.net (lgucdn.solbox-kaistcdni.net)|211.38.137.37|:80... connected.
HTTP request sent, awaiting response... 302 Found
Location: http://sbcdn.skt-kaistcdni.net/solbox-kaistcdni.net/?ucdn=100:0&dcdn=200:0&dcdn=400:0&dcdn=100:0&dcdn=200:0&dcdn=400:0&dcdn=100:0&hopcnt=3 [following]
--2012-11-03 19:58:38--  http://sbcdn.skt-kaistcdni.net/solbox-kaistcdni.net/?ucdn=100:0&dcdn=200:0&dcdn=400:0&dcdn=100:0&dcdn=200:0&dcdn=400:0&dcdn=100:0&hopcnt=3
Connecting to sbcdn.skt-kaistcdni.net (sbcdn.skt-kaistcdni.net)|1.237.57.3|:80... connected.
HTTP request sent, awaiting response... 302 Found
Location: http://sktcdn.lgu-kaistcdni.net/solbox-kaistcdni.net/?ucdn=100:0&dcdn=200:0&dcdn=400:0&dcdn=100:0&dcdn=200:0&dcdn=400:0&dcdn=100:0&dcdn=200:0&hopcnt=2 [following]
--2012-11-03 19:58:38--  http://sktcdn.lgu-kaistcdni.net/solbox-kaistcdni.net/?ucdn=100:0&dcdn=200:0&dcdn=400:0&dcdn=100:0&dcdn=200:0&dcdn=400:0&dcdn=100:0&dcdn=200:0&hopcnt=2
Connecting to sktcdn.lgu-kaistcdni.net (sktcdn.lgu-kaistcdni.net)|211.115.90.78|:80... connected.
HTTP request sent, awaiting response... 302 Found
Location: http://lgucdn.solbox-kaistcdni.net/solbox-kaistcdni.net/?ucdn=100:0&dcdn=200:0&dcdn=400:0&dcdn=100:0&dcdn=200:0&dcdn=400:0&dcdn=100:0&dcdn=200:0&dcdn=400:0&hopcnt=1 [following]
--2012-11-03 19:58:38--  http://lgucdn.solbox-kaistcdni.net/solbox-kaistcdni.net/?ucdn=100:0&dcdn=200:0&dcdn=400:0&dcdn=100:0&dcdn=200:0&dcdn=400:0&dcdn=100:0&dcdn=200:0&dcdn=400:0&hopcnt=1
Connecting to lgucdn.solbox-kaistcdni.net (lgucdn.solbox-kaistcdni.net)|211.38.137.37|:80... connected.
HTTP request sent, awaiting response... 302 Found
Location: http://sbcdn.skt-kaistcdni.net/solbox-kaistcdni.net/?ucdn=100:0&dcdn=200:0&dcdn=400:0&dcdn=100:0&dcdn=200:0&dcdn=400:0&dcdn=100:0&dcdn=200:0&dcdn=400:0&dcdn=100:0&hopcnt=0 [following]
--2012-11-03 19:58:38--  http://sbcdn.skt-kaistcdni.net/solbox-kaistcdni.net/?ucdn=100:0&dcdn=200:0&dcdn=400:0&dcdn=100:0&dcdn=200:0&dcdn=400:0&dcdn=100:0&dcdn=200:0&dcdn=400:0&dcdn=100:0&hopcnt=0
Connecting to sbcdn.skt-kaistcdni.net (sbcdn.skt-kaistcdni.net)|1.237.57.3|:80... connected.
HTTP request sent, awaiting response... 500 Internal Server Error
2012-11-03 19:58:38 ERROR 500: Internal Server Error.
```

Summary & Next Step

- Minor updates were made based on the comments to 00 version
- Some experiments performed with initial results
- Further tests will be performed and reported in the next IETF
- Propose to merge loop prevention mechanism with request routing redirection draft, draft-he-cdni-routing-request-redirection-03
- Any comments or suggestions for improvements are invited

Chair's Questions & Answers

- Do you feel that the scheme is well understood (i.e. what has to be signaled, how to e ncode it, how to process it)?
  - Encoding is simple. Currently part of URI query string or CNAME. It can also be encoded in JSON or other encoding formats
  - Signaling is done currently as a part of HTTP or DNS but can be done by RRRI interface prot ocol, for example, RRRI request & response message in the he's draft
  - Processing is a Request Router's operation behavior which is a part of redirection decision making process. And it is also simple. It can cover both loop prevention and detection with associated post detection processes
  - The same scheme can be equally applicable for both iterative and recursive redirections
  - For scheme's feasibility, performance impact, we performed experiments with some initial res ults

- Are you clear that it does not have any impact at all on other interfaces?
  - Loop detection is optional requirement then Capabilities advertisement needs to specify it as a part of capabilities
  - In case of metadata, if operational metadata is specified, loop detection or prevention can be one example.
  - For logging, loop detection & prevention event can be part of logging processes
  - For control interface, not affected by trigger interface but not clear for other control aspects si nce they are defined yet