Hochschule für Angewandte Wissenschaften Hamburg
Hamburg University of Applied Sciences

Content Delivery Networks Interconnection (CDNI)

Markus Vahlenkamp

AW1 seminar paper

Fakultät Technik und Informatik
Studiendepartment Informatik

Faculty of Engineering and Computer Science
Department of Computer Science

Contents

List of Figures

1 Introduction

The amount of data transferred through the Internet is continuously increasing. It is not just since the advent of smart phones, tablets and home entertainment devices that such a huge amount of data needs delivery. Thus, centralised serverfarms reach their limits in terms of available bandwidth and provided low latency communication.
One solution to cope with these problems was introduced by Content Delivery Networks (CDNs) in the late 1990th.
A CDN is a network of incorporating devices used to efficiently deliver digital content to content consumers. They provide numerous benefits also beyond the capability of delivering large amounts of data. Content providers are able to reduce the delivery costs of cacheable parts of their content by offloading the burden of delivery to the CDN providers equipment. This results in reduced management costs as well as a pay-per-use cost model for the content providers.
Furthermore utilising CDNs improves the consumers quality of experience (QoE). Due to an reduced network distance between content and consumer the perceived delays and latencies are reduced since subsequent request are served by the CDN servers (surrogate servers) cache.
The Network Service Providers (NSPs) derive the benefit of a decreased utilisation of their transit links and maybe also their backbone networks. The intensity of this impact depends on the position of the surrogate servers. This measure is anyhow able to lower costs by reducing equipment requirements and through transit cost savings.
A recent trend of NSPs is to build their own CDNs to subsequent introduce value-added services on top of it. They start to build their own Video-On-Demand libraries for their customers and also resell the CDN service to content providers. By doing so, NSPs try to actively create new revenue streams and reduce their operational expenses to further increase their benefit.
NSPs know their network the best. They know where bottlenecks exist or where unused resources are available. For this reason they are able to build a very efficient delivery infrastructure within their network. They are able to steer the traffic in order that the underlying network is used in an enhanced way regarding customer perceived delay and bandwidth utilisation.
Nevertheless some difficulties remain for the NSPs that build there own CDNs and try to com-

pete with global CDN providers. Compared to them the NSPs do have a limited footprint, the geographic distribution it can offer. Their CDNs are optimised with the focus of delivering content to their own customer base, whereas global CDN providers situate their surrogate servers within different networks or peer with them to be able to provide considerably more customers with an increased QoE.
Thus the requirements of interconnecting these separated CDN islands to derive an increased advantage is desired.
Despite of this known potential for improvement, there is no standard method of interconnecting CDNs defined so far. Merely a couple of experimental CDN interconnection trails exist today which leads to the requirement of specifying standardised methods and interfaces for the interconnection of CDNs.

2 Content Delivery Networks

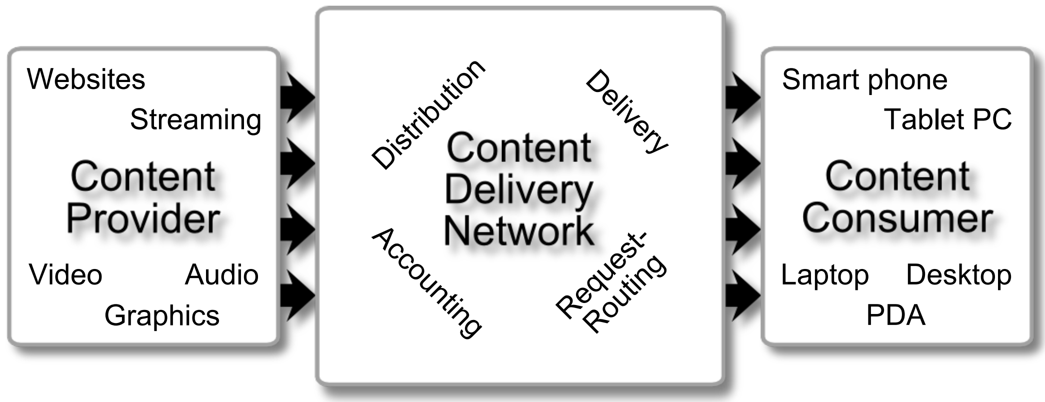

Figure 1: Conceptional CDN overview [Vahlenkamp 2011; Buyya u. a. 2006]

A CDN is a federation of computer systems aiming at an efficient delivery of digital content to content consumers.
Figure 1 depicts an abstract and high level overview of CDNs. As shown CDNs connect content providers with end users, the content consumers. The content provider makes its data, for instance pictures, video or audio-files, available for delivery to content consumers within the CDN.
Whenever a content consumer requests content, the geographically distributed surrogate servers of the CDN are utilised to efficiently deliver the requested data.
As depicted in figure 2 a CDN consists of subsystems that serve different purposes. The general separation leads to four subsystems, namely Request-Routing, Delivery, Distribution and Accounting System. They are explained in more detail in what follows.

Request-Routing System The Request-Routing System steers content requests to appropriate surrogate servers. To achieve this, it interacts with the Delivery System to retrieve status information about the currently cached content on the different surrogate servers.

Delivery System The Delivery System consists of surrogate servers that deliver the content to content consumers.

Distribution System The Distribution System is responsible for acquiring the content from the content providers origin servers and for its distribution within the CDN. This is achieved by pre-positioning it on the surrogate servers or by an on demand acquisition process for example.

Accounting System The Accounting System interacts with all the above mentioned subsystems. It collects and analysis the log files and statistics for the purpose of monitoring and billing the obtained services.

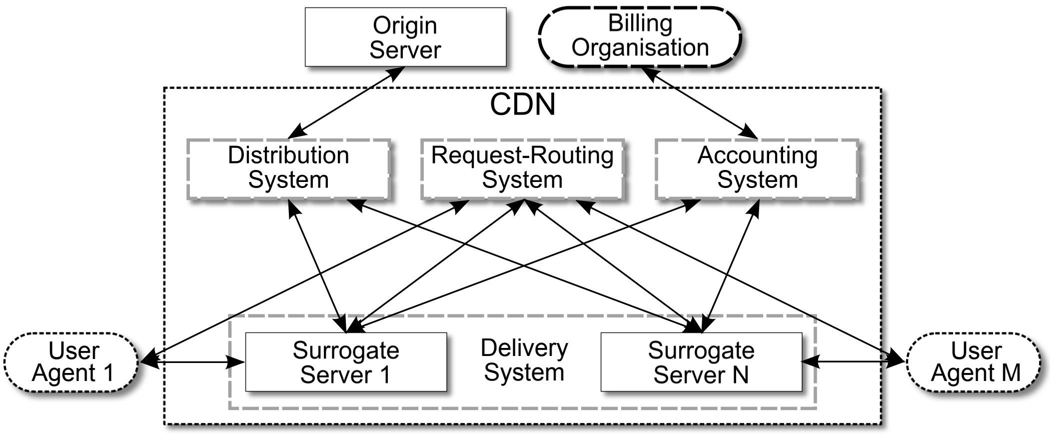

Figure 2: CDN components [Vahlenkamp 2011; Day u. a. 2003; Hofmann und Beaumont 2005; Yin u. a. 2010]

All of the interfaces between the CDN subsystems shown in figure 2 are vendor specific. Therefore it is impossible to mix CDN components of different vendors. Furthermore it is desired to keep CDNs of different authorities separate. Thus a specialised and standardised inter-CDN interface is required for the purpose of connecting different CDNs.

3 CDNI

The core motivation for interconnecting CDNs is to connect independent CDN islands with the goal of making them deliver content on behalf of each other [Bertrand u. a. 2011a].
NSPs are for instance not able to provide the same footprint that global CDN providers are able to provide [Pathan 2009]. Hence content providers are forced to contract with various NSPs to gain an equal service when using NSP operated CDNs [Eurescom 2010].
By connecting their CDNs the NSPs can overcome this shortcoming and rise there attraction to content providers.

3.1 Use Cases

The strongest argument for interconnecting CDNs is to extend their geographical footprint. This could be an interconnection between CDNs of different CDN provider or the interconnection of different vendors CDN products operated by the same CDN provider.
Gaining a larger footprint is only one capability the interconnection of CDNs is able to provide [Bertrand u. a. 2011a]. Furthermore the CDNs could utilise one another to provide a broader range of technological capabilities to content providers and content consumers. This could for instance be the capability of transcoding video content which is increasingly useful since the variate of user devices is growing. Devices like laptops, smart phones and tablets have different capabilities and requirements with respect to their CPU power, available bandwidth and of course their screen size. Thus optimizing the video footage for the end device results in an improved QoE.
The interconnectability would also enable a CDN provider to mix the CDN products of different vendors on its own network and let them work as one united feature rich CDN.
Another reason is the event of a flash crowd [Arlitt und Jin 2000] or when there is maintenance work going on in one CDN, the user requests could be offload to an adjacent CDN, one with available capacity reserves.
All these measures can rise the attraction of the interconnected CDNs, because they will operate cooperatively, extending the footprint, the available technologies as well aslowering the management costs.

3.2 Operation

The interconnected CDNs will cooperate to provide consumers with content. Figure 3 illustrates a scenario where two CDNs are interconnected and thus are able to cooperate when delivering the content providers data to content consumers.

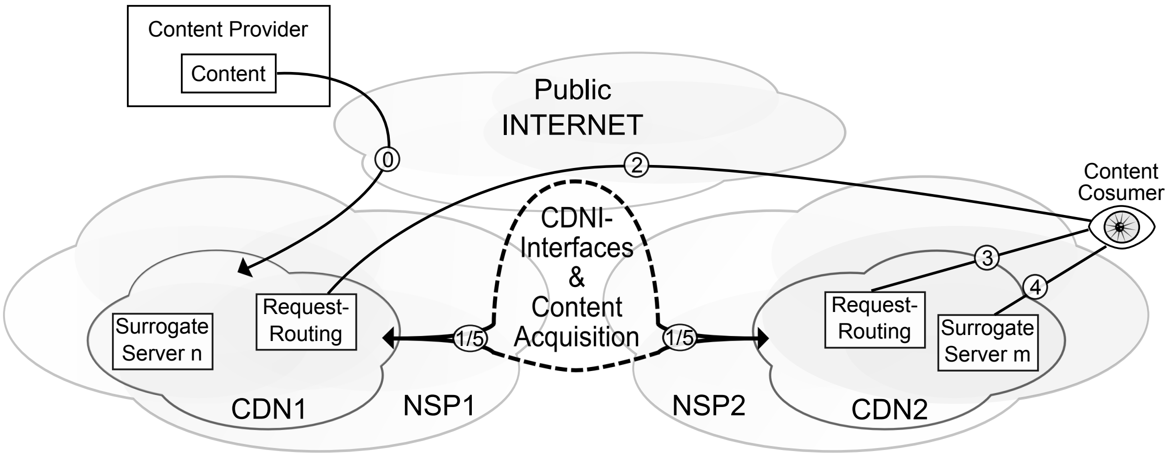

Figure 3: Simplified view of the CDNI interface when leveraged to delivering content.

It is assumed that the CDN interconnection is already in place. Thus no further initialisation of the Content Delivery Network Interconnection (CDNI) interface is required.
In the illustrated scenario CDN1 serves as the authoritative upstream CDN whereas CDN2 takes the role of the downstream CDN. This is CDN2 acts on behalf of CDN1 which again acts on behalf of the content provider.
The actual content consumer is located next to CDN2. Because of its proximity the content consumer is served best by CDN2.
The following steps describe how the CDNI interface is leveraged to deliver the content to the content consumer. Starting at the point where the content consumer requests the content up till the particular content delivery.

0. As a pre-step the content provider places its content within the authoritative CDN1.
1. The upstream CDN has to keep track of the downstream CDN, it has to take into account which clients the downstream CDN is willing or able to serve. Furthermore it must determine whether CDN2 is able to handle requests or if it is maybe overloaded. In the latter case the upstream CDN would temporarily avoid utilising CDN2 for content delivery.
2. The content consumer requests the content. The request is processed by the Request-Routing System of the authoritative CDN1, which in turn decides that CDN2 should

serve the request. The content consumer is redirected to the Request-Routing System of the downstream CDN2.

3. The Request-Routing System of CDN2 determines the surrogate server to process the content consumers request and redirects it to that server via its response message.

4. The content consumers user agent connects to the selected surrogate server to finally request and retrieve the desired content.

5. For the first request of a particular piece of content destined to the surrogate server of the downstream CDN, it has to request the content from its upstream CDN. Through the Delivery System the surrogate server determines where to acquire the content from. In this case the content is acquired directly form the authoritative CDN, delivered to the content consumer and at the same time cached for subsequent requests.
   This process of acquiring the content also applies if the requested content is no more cached on the surrogate server.

When taking a closer look at the interconnection of CDNs, one can see that it results in interconnecting the subsystems of the CDNs. The subsystems have to communicate, they have to exchange informations to cooperate.
Figure 4 shows two interconnected CDNs in more detail compared to Figure 3. The interfaces between the adjacent CDN Control, Logging, Request-Routing and Distribution Systems are subject of the CDNI efforts, whereas the interfaces for content acquisition, the method used to request as well as to deliver the content, are out of scope of the Internet Engineering Task Force (IETF) CDNI working group standardisation efforts.

3.3 Interface Requirements

There are several requirements that CDNI has to take into account [Bertrand u. a. 2011a; Louedec u. a. 2011]. It needs to enable the pre-positioning of content or metadata within the downstream CDN in order to lower the content delivery delay and to provide the service level requested by the content provider. On the other hand content must also be acquirable on demand, so that only the information about where to locate the actual content if needed is available within the adjacent CDN.
CDNs also need to provide methods to assure cache coherence. For instance the purge operations used to invalidate content also needs to reach the downstream CDNs, so that the content is no longer available for delivery to content consumers.
Digital rights management needs to be taken into account. For this Geo-Blocking functionalities must be part of the CDNI specification. Through Geo-Blocking geographical regions with their corresponding IP-Ranges are prevented from retrieving particular content.

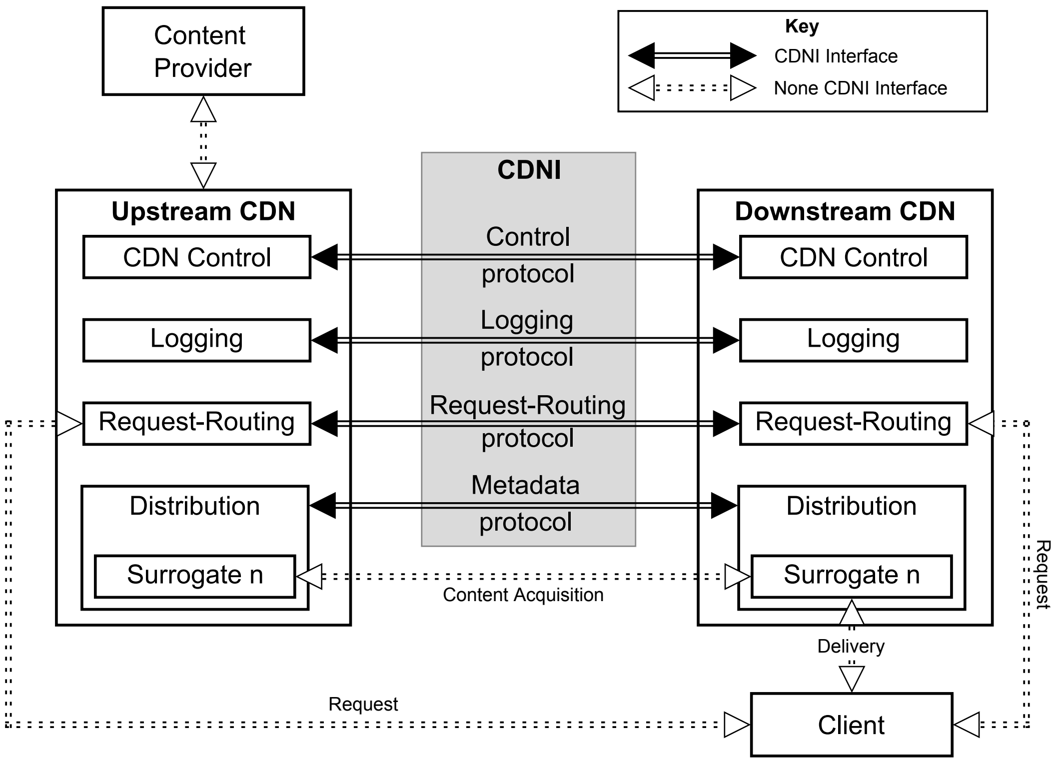

Figure 4: CDNI interfaces [Niven-Jenkins u. a. 2011]

Through this functionality CDN providers can for instance enforce that some content is only available within a particular country.
Another restriction method required by CDNs and for this also by CDNI is the timely restriction of the content availability. A common technique in use to prevent children from consuming content that doesn't fit their age. Content providers assign time restrictions to critical content, so that it is only available at night-times and for this should not be available to children.

The following provides a collection of interface-specific requirements that exist for CDNI. It represents an abstract of the requirements listed in [Niven-Jenkins u. a. 2011].

Control Interface The CDNI Control Interface must allow the Control Systems of interconnected CDNs to exchange information to establish, update or terminate the CDN interconnection. Also the bootstrapping and configuration of the other interfaces is part of the CDNI Control Interfaces responsibilities.

CDNI Logging Interface The CDNI Logging Interface enables the interconnected CDNs to exchange information about the particular data that was transmitted as well as the amount. All this information may then be used for billing purpose and as a data source for analytical and monitoring systems. For all of these purposes, information like IP addresses and the user agent (e.g. browser) used to request the data are also required.

CDNI Request-Routing Interface The CDNI Request-Routing Interface is used to exchange information about the Request-Routing Systems between upstream and downstream CDNs. The upstream CDN must be able to request information about the downstream CDN's ability to serve the requested content, as well as the redirection information that should be passed to the content consumer's user agent.
The CDNI Request-Routing Interface must also be able to cope with various protocols that are used to redirect user agents, in particular (but not exclusively) DNS [Mockapetris 1987], HTTP [Fielding u. a. 1999] and RTSP [Schulzrinne u. a. 1998]. The CDNI Request-Routing Interface should end up as a request/response interface that leverage's existing request/response protocols. It is not intended to invent a new request/response protocol, furthermore protocols like XML-RPC, HTTP query should be considered as candidates. Thus only the syntax and semantic of the redirection request responses have to be specified as well as some standard behaviour, e.g. how to respond to malformed requests or responds.
Also the previously mentioned information about footprint, available resources and load need to be provided by the CDNI Request-Routing Interface as well.

CDNI Metadata Interface The CDNI metadata Interface must enable the Distribution System in a downstream CDN to obtain CDNI metadata from the upstream CDN in order to properly deal with:

- content requests from user agents
- redirect requests received from authoritative CDNs

By using the CDNI Metadata Interface the upstream CDN must be enabled to distribute, update and remove metadata concerning its content to the downstream CDNs.
The CDNI Metadata Interface will probably also end up being a request/response interface and should as well as the CDNI Request-Routing Interface leverage existing request/response protocols. Thus also only the syntax, semantic and standard behaviour needs specification.

The design of these CDNI interface must also provide the ability for enhancements. For instance adaptive video streaming techniques like the recent Moving Picture Experts Group (MPEG) Dynamic Adaptive Streaming over HTTP (DASH) standard [Richardson 2003].
Adaptive video streaming is a technique to increase the QoE when consuming video content and is thus a functionality desirably supported by CDNI.
The video content is available in different bit rates. According to the available bandwidth the video file with the highest bit rate that does not exceed the available bandwidth[1] is chosen

[1] Also parameters like CPU performance, battery level etc. could be taken into account.

and delivered to the content consumer. The resulting play-out video can be a mixture of different bit rates. Through adaptive video streaming the content consumer should always experience the highest available quality of the content, without getting interrupted because of insufficient bandwidth available for the particular bit rate to be transmitted without dropouts.
These techniques and functionalities, as well as future developments have to be supported at best. Preferably out-of-the-box or by protocol extensions.

4 Related work

4.1 Early IETF work: CDI

Within the time frame from 2001 to 2003 the IETF CDI working group was active. The working group target was to standardise some early CDN interconnection.
In 2003 the working group was closed because of the consolidation of the CDN market, resulting in a flattening interest in the CDI working group.
The achievements of this working group have been three informational Request for Comment (RFC) documents.

RFC3466 The document "A model for content distribution internetworking" [Day u. a. 2003] introduces content networks and content internetworking and defines a common vocabulary for the CDI working group.

RFC3568 The document "Known Content Network (CN) Request-Routing Mechanisms" [Barbir u. a. 2003] summarises request-routing mechanisms categorised by DNS request-routing, transport-layer request-routing and application-layer request-routing.

RFC3570 The document "Content Internetworking (CDI) Scenarios" [Rzewski u. a. 2003] provides examples of scenarios that may occur when two content networks decide to interconnect. The scenarios presented seek to provide some concrete examples of what content internetworking is and also to provide a basis for evaluating content internetworking proposals.

4.2 Recent CDNI Experiment

France Telecom-Orange Labs set up an experiment [Bertrand u. a. 2011b] with CDN equipment of two different vendors. The goal of the experiment was to study and analyse the (minimum) requirements for interconnecting two CDNs.
For their tests France Telecom-Orange Labs used two different vendors (further referenced by CDN A and CDN B) CDN solutions. Both CDNs covered a different country. Thus content consumers from Country A should be served by CDN A, content consumers from Country B by CDN B.
The result of the experiment was, like expected, that the internally used protocols of the CDNs are proprietary and thus not compatible with each other.
Nevertheless it was possible to interconnect these CDNs. Some essential information exchange was achieved through "in-band" communication, the encoding of information within the Uniform Resource Identifiers (URIs) used to locate the content.
The overall insight however was, that a proper interconnection of CDNs has to also utilise some sort of an "out-of-band" communication to be able to make use of the full potential of an interconnected CDN solution. The outcomes of this experiment mostly resulted in contribution to the CDNI working groups problem statement draft [Niven-Jenkins u. a. 2011].

5 Conclusion

This paper presented a brief overview of CDNs and introduced the topic of interconnecting them for the purpose of a cooperated content delivery infrastructure.
We represented arguments how CDNI could help raising the attraction of NSP-driven CDNs and thus how they will be able to increase their economic incentives.
Throughout this paper we highlighted that many requirements need to be taken into account when designing the CDNI interface. For instance the independence of particular request-routing or content acquisition methods.
Through all this the NSP-driven CDNs will be able to operate cooperatively. They will extend their footprint, the available technologies and at the same time lower the management costs. All this will lead to a rising attraction of these interconnected NSP-driven CDNs.

The particular CDNI working group was formed only in April 2011. Hence many things still need to be figured out and a lot further research effort is required to standardise a fully functioning CDN interconnection.

Glossary

A

authoritative CDN The CDN with an direct relationship with the Content provider for distribution and delivery of that Content providers content.

C

CDI Content Distribution Internetworking.

CDN *see* Content Delivery Network

CDN provider A service provider that operates a CDN. The content delivery service is typically used by Content providers or other CDN Providers for content delivery.

CDNI *see* Content Delivery Network Interconnection

content consumer The end user of the system who uses the User agent to retrieve and consume content.

Content Delivery Network Network infrastructure in which the network elements cooperate at layers 4 through layer 7 for an improved effectiveness while delivering content to User agents.

Content Delivery Network Interconnection A relationship between CDNs within which a CDN (Downstream CDN) provides content delivery services to Content consumers on behalf of another CDN (Upstream CDN).

content provider Provides content services to Content consumers. The CSP may own content or may license distributions rights from another party.

D

DASH Dynamic Adaptive Streaming over HTTP.

downstream CDN For a given End User request, the CDN to which the request is redirected by the other CDN (Upstream CDN).

F

flash crowd A spontaneous and temporary increase in friendly (no DoS-Attack) content requests caused by a hype with regard to a particular content.

I

IETF Internet Engineering Task Force.

M

MPEG Moving Picture Experts Group.

N

Network Service Provider Provider of network based services / connectivity.

NSP *see* Network Service Provider

O

origin server A content server operated by the Content provider where the Surrogate servers of the Authoritative CDN may request the content that is to be delivered.

Q

QoE *see* quality of experience

QoS *see* quality of service

quality of experience The service quality that is subjectively perceived by a Content consumer. In contrast to Quality of Service (QoS) the QoE is not measurable because it relies on subjective impressions.

quality of service The measurable quality of a provided service in general with respect to delay, jitter and throughput.

R

RFC Request for Comment.

S

surrogate server A server that interacts with other components of the CDN to control the flow of content and also delivers content to the User agents.

U

upstream CDN For a given End User request, the CDN that redirects the request to the other CDN.

URI Uniform Resource Identifier.

user agent The software that is used by the Content consumer to communicate with the CDNs. This could also include set-top-boxes, browsers, mediaplayer etc. and is not limited to the HTTP protocol.

Bibliography

[Eurescom 2010] P1955 - CDN interconnection / Eurescom. August 2010. – D3 - Summary of study results

[Arlitt und Jin 2000] ARLITT, M. ; JIN, T.: A workload characterization study of the 1998 World Cup Web site. In: *Network, IEEE* 14 (2000), Mai, Nr. 3, S. 30–37. – ISSN 0890-8044

[Barbir u. a. 2003] BARBIR, A. ; CAIN, B. ; NAIR, R. ; SPATSCHECK, O.: Known Content Network (CN) Request-Routing Mechanisms / RFC-Editor. URL http://www.rfc-editor.org/rfc/rfc3568.txt, Juli 2003 (3568). – RFC

[Bertrand u. a. 2011a] BERTRAND, Gilles ; EMILE, Stephan ; WATSON, Grant ; BURBRIDGE, Trevor ; EARDLEY, Philip ; MA, Kevin: Use Cases for Content Delivery Network Interconnection / RFC-Editor. URL http://tools.ietf.org/html/draft-ietf-cdni-use-cases-00, September 2011 (draft-ietf-cdni-use-cases-00). – Internet-Draft

[Bertrand u. a. 2011b] BERTRAND, Gilles ; FAUCHEUR, Francois L. ; PETERSON, Larry: Content Distribution Network Interconnection (CDNI) Experiments / RFC-Editor. Februar 2011 (draft-bertrand-cdni-experiments-02). – Internet-Draft

[Buyya u. a. 2006] BUYYA, Rajkumar ; PATHAN, Al-Mukaddim K. ; BROBERG, James ; TARI, Zahir: A Case for Peering of Content Delivery Networks. In: *Distributed Systems Online, IEEE* 7 (2006), Oktober, Nr. 10, S. 3. – ISSN 1541-4922

[Day u. a. 2003] DAY, M. ; CAIN, B. ; TOMLINSON, G. ; RZEWSKI, P.: A Model for Content Internetworking (CDI) / RFC-Editor. URL http://www.rfc-editor.org/rfc/rfc3466.txt, Februar 2003 (3466). – RFC

[Fielding u. a. 1999] FIELDING, R. ; GETTYS, J. ; MOGUL, J. ; FRYSTYK, H. ; MASINTER, L. ; LEACH, P. ; BERNERS-LEE, T.: Hypertext Transfer Protocol – HTTP/1.1 / RFC-Editor. URL http://www.rfc-editor.org/rfc/rfc2616.txt, Juni 1999 (2616). – RFC. Updated by RFCs 2817, 5785

[Hofmann und Beaumont 2005] HOFMANN, Markus ; BEAUMONT, Leland R.: *Content Networking: Architecture, Protocols, and Practice (The Morgan Kaufmann Series in Networking)*. San Francisco, CA, USA : Morgan Kaufmann Publishers Inc., 2005. – ISBN 1-55860-834-6

[Louedec u. a. 2011] LOUEDEC, Yannick L. ; TIMMERER, Christian ; SPIROU, Spiros ; GRIFFIN, David ; ZAHARIADIS, Theodore: Catalogue of Advanced Use Cases for Content Delivery Network Interconnection / RFC-Editor. Oktober 2011 (draft-fmn-cdni-advanced-use-cases-00). – Internet-Draft

[Mockapetris 1987] MOCKAPETRIS, P.V.: Domain names - concepts and facilities / RFC-Editor. URL `http://www.rfc-editor.org/rfc/rfc1034.txt`, November 1987 (1034). – RFC. Updated by RFCs 1101, 1183, 1348, 1876, 1982, 2065, 2181, 2308, 2535, 4033, 4034, 4035, 4343, 4035, 4592, 5936

[Niven-Jenkins u. a. 2011] NIVEN-JENKINS, Ben ; FAUCHEUR, Francois L. ; BITAR, Nabil: Content Distribution Network Interconnection (CDNI) Problem Statement / RFC-Editor. URL `http://tools.ietf.org/html/draft-ietf-cdni-problem-statement-01`, Oktober 2011 (draft-ietf-cdni-problem-statement-01). – Internet Draft

[Pathan 2009] PATHAN, Mukaddim: *Utility-Oriented Internetworking of Content Delivery Networks*, Department of Computer Science and Software Engineering, The University of Melbourne, Dissertation, November 2009. – URL `http://www.cloudbus.org/students/MukaddimPhDThesis2009.pdf`

[Richardson 2003] RICHARDSON, Iain E. G.: *H.264 and MPEG-4 Video Compression: Video Coding for Next-generation Multimedia*. New York, NY, USA : John Wiley & Sons, Inc., 2003

[Rzewski u. a. 2003] RZEWSKI, P. ; DAY, M. ; GILLETTI, D.: Content Internetworking (CDI) Scenarios / RFC-Editor. URL `http://www.rfc-editor.org/rfc/rfc3570.txt`, Juli 2003 (3570). – RFC

[Schulzrinne u. a. 1998] SCHULZRINNE, H. ; RAO, A. ; LANPHIER, R.: Real Time Streaming Protocol (RTSP) / RFC-Editor. URL `http://www.rfc-editor.org/rfc/rfc2326.txt`, April 1998 (2326). – RFC

[Vahlenkamp 2011] VAHLENKAMP, Markus: *Content Delivery Networks - Chancen und Konzepte für Internet Service Provider*, Hochschule für Angewandte Wissenschaften Hamburg, Bachelor-Thesis, April 2011

[Yin u. a. 2010] YIN, Hao ; LIU, Xuening ; MIN, Geyong ; LIN, Chuang: Content delivery networks: A bridge between emerging applications and future IP networks. In: *Network, IEEE* 24 (2010), Juli, Nr. 4, S. 52–56. – ISSN 0890-8044