

DECLARATION

I, Alexa Morris, based on my personal knowledge and information, hereby declare as follows:

1. I am Managing Director of the IETF Administration LLC and have held that position since the LLC was formed in August 2018. Prior to that, starting on January 1, 2008, I was the Executive Director of the Internet Engineering Task Force, which was an activity of the Internet Society. Since the business of IETF did not change in any materially relevant manner with the formation of the LLC, I will collectively refer to both the activity and the LLC as IETF.

2. One of my responsibilities with IETF has been to act as the custodian of Internet-Drafts and records relating to Internet-Drafts. I am familiar with the record keeping practices relating to Internet-Drafts, including the creation and maintenance of such records.

3. I hereby declare that all statements made herein are of my own knowledge and information contained in the business records of IETF and are true, and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements may be punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code.

4. If depositions regarding the information in this declaration are required, the deposition should be taken by phone or videoconference or, if it must be in person, should be in California.

5. Since 1998, it has been the regular practice of the IETF to publish Internet-Drafts and make them available to the public on its website at www.ietf.org (the IETF website). The IETF maintains copies of Internet-Drafts in the ordinary course of its regularly conducted activities.

6. Any Internet-Draft published on the IETF website was reasonably accessible to the public and was disseminated or otherwise available to the extent that persons interested and ordinarily skilled in the subject matter or art exercising reasonable diligence could have located it. In particular, the Internet-Drafts were indexed and searchable on the IETF website.

7. Internet-Drafts are posted to an IETF online directory. When an Internet-Draft is published, an announcement of its publication that describes the Internet-Draft is disseminated. Typically, that dated announcement is made within 24 hours of the publication of the Internet-Draft. The announcement is kept in the IETF email archive and the date is affixed automatically.

8. The records of posting the Internet-Drafts in the IETF online repository are kept in the course of the IETF's regularly conducted activity and ordinary course of business. The records are made pursuant to established procedures and are relied upon by the IETF in the performance of its functions.

9. It is the regular practice of the IETF to make and keep the records in the online repository.


10. Exhibit 1 is a true and correct copy of an announcement of the publication of draft-ietf-cdi-model-00.txt, titled "A Model for Content Internetworking (CDI)." I have determined that an announcement of the publication of this Internet-Draft was made on February 22, 2002. Therefore, based on the normal practice of the IETF, that Internet-Draft was reasonably available to the public within 24 hours of that announcement. At that time, the Internet-Draft would have been disseminated or otherwise available to the extent that persons interested and ordinarily skilled in the subject matter or art, exercising reasonable diligence, could have located it.

11. Exhibit 2 is a true and correct copy of an announcement of the publication of draft-green-cdn-p-gen-arch-03.txt, titled "Content Internetworking Architectural Overview." I have determined that an announcement of the publication of this Internet-Draft was made on March 2, 2001. Therefore, based on the normal practice of the IETF, that Internet-Draft was reasonably available to the public within 24 hours of that announcement. At that time, the Internet-Draft would have been disseminated or otherwise available to the extent that persons interested and ordinarily skilled in the subject matter or art, exercising reasonable diligence, could have located it.

12. Exhibit 3 is a true and correct copy of an announcement of the publication of draft-amini-cdi-distribution-reqs-02.txt, titled "Distribution Requirements for Content Internetworking." I have determined that an announcement of the publication of this Internet-Draft was made in March 2001. Therefore, based on the normal practice of the IETF, that Internet-Draft was reasonably available to the public within 24 hours of that announcement. At that time, the Internet-Draft would have been disseminated or otherwise available to the extent that persons interested and ordinarily skilled in the subject matter or art, exercising reasonable diligence, could have located it.

Pursuant to Section 1746 of Title 28 of United States Code, I declare under penalty of perjury under the laws of the United States of America that the foregoing is true and correct and that the foregoing is based upon personal knowledge and information and is believed to be true.

Date: August 23, 2002

By: 
Alexa Morris

4884-7542-0207

EXHIBIT 1

A Model for Content Internetworking (CDI)
draft-ietf-cdi-model-00.txt

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of RFC2026.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on August 23, 2002.

Copyright Notice

Copyright (C) The Internet Society (2002). All Rights Reserved.

Abstract

Content [distribution] internetworking (CDI) is the technology for interconnecting content networks, sometimes previously called "content peering" or "CDN peering." A common vocabulary helps the process of discussing such interconnection and interoperation. This document introduces content networks and content internetworking, and defines elements for such a common vocabulary.

Table of Contents

1.	Introduction	3
2.	Content Networks	4
2.1	Problem Description	4
2.2	Caching Proxies	5
2.3	Server Farms	6
2.4	Content Distribution Networks	7
2.4.1	Historic Evolution of CDNs	9
2.4.2	Describing CDN Value: Reach and Scale	9
3.	Content Network Model Terms	11
4.	Content Internetworking	14
5.	Content Internetworking Model Terms	15
6.	Security Considerations	18
7.	Acknowledgements	19
	References	20
	Authors' Addresses	21
	Full Copyright Statement	22

1. Introduction

Content networks are of increasing importance to the overall architecture of the Web. This document presents a vocabulary for use in developing technology for interconnecting content networks, or "content internetworking."

The accepted name for the technology of interconnecting content networks is "content internetworking." For historical reasons, we abbreviate this term using the acronym CDI (from "content distribution internetworking"). Earlier names relied on analogy with peering and interconnection of IP networks; thus we had "content peering" and "CDN peering". All of these other names are now deprecated, and we have worked to establish consistent usage of "content internetworking" and "CDI" throughout the drafts of the IETF CDI group.

The terminology in this document builds from the previous taxonomy of web caching and replication in RFC 3040 [3]. In particular, we have attempted to avoid the use of the common terms "proxies" or "caches" in favor of more specific terms defined by that document, such as "caching proxy."

Section 2 provides background on content networks. Section 3 introduces the terms used for elements of a content network and explains how those terms are used. Section 4 provides additional background on interconnecting content networks, following which Section 5 introduces additional terms and explains how those internetworking terms are used.

[Note to RFC Editor: This entire paragraph may be deleted so as to avoid references to internet-drafts in RFCs.] The IETF CDI effort has produced a number of other documents related to content internetworking. Other documents providing general information about CDI are: "Content Internetworking Scenarios" [5], which enumerates scenarios for content-internetworking-related interactions; "Content Internetworking Architectural Overview" [4], which gives an overall architecture of the elements for CDI; and "Known CDN Request-Routing Mechanisms" [7], which summarizes known mechanisms for request-routing. In addition, there are documents describing the requirements for various aspects of CDI: "Request-Routing Requirements for Content Internetworking" [8], "Distribution Requirements for Content Internetworking" [9], and "Content Internetworking (CDI) Authentication, Authorization, and Accounting Requirements" [6]

2. Content Networks

The past several years have seen the evolution of technologies centered around "content." Protocols, appliances, and entire markets have been created exclusively for the location, download, and usage tracking of content. Some sample technologies in this area have included web caching proxies, content management tools, intelligent "web switches", and advanced log analysis tools.

When used together, these tools form new types of networks, dubbed "content networks". Whereas network infrastructures have traditionally processed information at layers 1 through 3 of the OSI stack, content networks include network infrastructure that exists in layers 4 through 7. Whereas lower-layer network infrastructures centered on the routing, forwarding, and switching of frames and packets, content networks deal with the routing and forwarding of requests and responses for content. The units of transported data in content networks, such as images, movies, or songs, are often very large and may span hundreds or thousands of packets.

Alternately, content networks can be seen as a new virtual overlay to the OSI stack: a "content layer", to enable richer services that rely

on underlying elements from all 7 layers of the stack. Whereas traditional applications, such as file transfer (FTP), relied on underlying protocols such as TCP/IP for transport, overlay services in content networks rely on layer 7 protocols such as HTTP or RTSP for transport.

The proliferation of content networks and content networking capabilities gives rise to interest in interconnecting content networks and finding ways for distinct content networks to cooperate for better overall service.

2.1 Problem Description

Content networks typically play some role in solving the "content distribution problem". Abstractly, the goal in solving this problem is to arrange a rendezvous between a content source at an origin server and a content sink at a viewer's user agent. In the trivial case, the rendezvous mechanism is that every user agent sends every request directly to the origin server named in the host part of the URL identifying the content.

As the audience for the content source grows, so do the demands on the origin server. There are a variety of ways in which the trivial system can be modified for better performance. The apparent single logical server may in fact be implemented as a large "farm" of server machines behind a switch. Both caching proxies and reverse caching

Day, et al. Expires August 23, 2002 [Page 4]
Internet-Draft CDI Model February 2002

proxies can be deployed between the client and server, so that requests can be satisfied by some cache instead of by the server.

For the sake of background, several sample content networks are described in the following sections that each attempt to address this problem.

2.2 Caching Proxies

A type of content network that has been in use for several years is a caching proxy deployment. Such a network might typically be employed by an ISP for the benefit of users accessing the Internet, such as through dial or cable modem.


In the interest of improving performance and reducing bandwidth utilization, caching proxies are deployed close to the users. These users are encouraged to send their web requests through the caches rather than directly to origin servers, such as by configuring their browsers to do so. When this configuration is properly done, the user's entire browsing session goes through a specific caching proxy. That caching proxy will therefore contain the "hot set" of all Internet content being viewed by all of the users of that caching proxy.

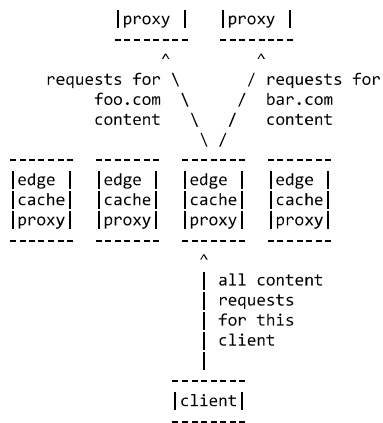
When a request is being handled at a caching proxy on behalf of a user, other decisions may be made, such as:

- o A provider that deploys caches in many geographically diverse locations may also deploy regional parent caches to further aggregate user requests and responses. This may provide additional performance improvement and bandwidth savings. When parents are included, this is known as hierarchical caching.
- o Using rich parenting protocols, redundant parents may be deployed such that a failure in a primary parent is detected and a backup is used instead.
- o Using similar parenting protocols, requests may be partitioned such that requests for certain content domains are sent to a specific primary parent. This can help to maximize the efficient use of caching proxy resources.

The following diagram depicts a hierarchical cache deployment as described above:

Day, et al. Expires August 23, 2002 [Page 5]
Internet-Draft CDI Model February 2002





Note that this diagram shows only one possible configuration, but many others are also useful. In particular, the client may be able to communicate directly with multiple caching proxies. RFC 3040 [3] contains additional examples of how multiple caching proxies may be used.

2.3 Server Farms


Another type of content network that has been in widespread use for several years is a server farm. A typical server farm makes use of a so-called "intelligent" or "content" switch (i.e. one that uses information in OSI layers 4-7). The switch examines content requests and dispatches them among a (potentially large) group of servers.

Some of the goals of a server farm include:

- o Creating the impression that the group of servers is actually a single origin site.

- o Load-balancing of requests across all servers in the group.
- o Automatic routing of requests away from servers that fail.
- o Routing all requests for a particular user agent's session to the same server, in order to preserve session state.

The following diagram depicts a simple server farm deployment:



A similar style of content network (that is, deployed close to servers) may be constructed with surrogates [3] instead of a switch.

2.4 Content Distribution Networks

Both hierarchical caching and server farms are useful techniques, but have limits. Server farms can improve the scalability of the origin server. However, since the multiple servers and other elements are typically deployed near the origin server, they do little to improve performance problems that are due to network congestion. Caching proxies can improve performance problems due to network congestion (since they are situated near the clients) but they cache objects based on client demand. Caching based on client demand performs poorly if the requests for a given object, while numerous in aggregate, are spread thinly among many different caching proxies. (In the worst case, an object could be requested n times via n distinct caching proxies, causing n distinct requests to the origin server -- or exactly the same behavior that would occur without any caching proxies in place.)

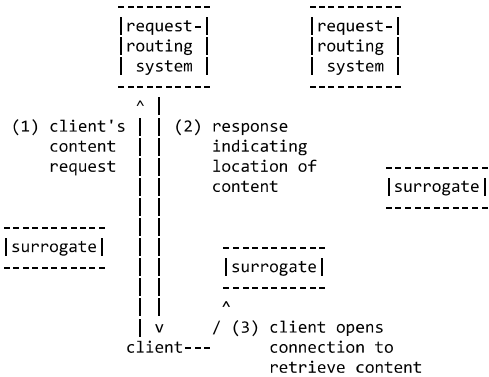
Thus, a content provider with a popular content source can find that it has to invest in large server farms, load balancing, and high-bandwidth connections to keep up with demand. Even with those investments, the user experience may still be relatively poor due to congestion in the network as a whole.

To address these limitations, another type of content network that has been deployed in increasing numbers in recent years is the CDN (Content Distribution Network or Content Delivery Network). A CDN essentially moves server-farm-like configurations out into network locations more typically occupied by caching proxies. A CDN has multiple replicas of each content item being hosted. A request from a browser for a single content item is directed to a "good" replica, where "good" usually means that the item is served to the client quickly compared to the time it would take fetch it from the origin server, with appropriate integrity and consistency. Static information about geographic locations and network connectivity is usually not sufficient to do a good job of choosing a replica. Instead, a CDN typically incorporates dynamic information about network conditions and load on the replicas, directing requests so as to balance the load.

Compared to using servers and surrogates in a single data center, a CDN is a relatively complex system encompassing multiple points of presence, in locations that may be geographically far apart. Operating a CDN is not easy for a content provider, since a content provider wants to focus its resources on developing high-value content, not on managing network infrastructure. Instead, a more typical arrangement is that a network service provider builds and operates a CDN, offering a content distribution service to a number of content providers.

A CDN enables a service provider to act on behalf of the content provider to deliver copies of origin server content to clients from multiple diverse locations. The increase in number and diversity of location is intended to improve download times and thus improve the user experience. A CDN has some combination of a content-delivery infrastructure, a request-routing infrastructure, a distribution infrastructure, and an accounting infrastructure. The content-delivery infrastructure consists of a set of "surrogate" servers [3] that deliver copies of content to sets of users. The request-routing infrastructure consists of mechanisms that move a client toward a rendezvous with a surrogate. The distribution infrastructure consists of mechanisms that move content from the origin server to the surrogates. Finally, the accounting infrastructure tracks and collects data on request-routing, distribution, and delivery functions within the CDN.

The following diagram depicts a simple CDN as described above:



2.4.1 Historic Evolution of CDNs

The first important use of CDNs was for the distribution of heavily-requested graphic files (such as GIF files on the home pages of popular servers). However, both in principle and increasingly in practice, a CDN can support the delivery of any digital content -- including various forms of streaming media. For a streaming media CDN (or media distribution network or MDN), the surrogates may be operating as splitters (serving out multiple copies of a stream). The splitter function may be instead of, or in addition to, a role as

a caching proxy. However, the basic elements defined in this model are still intended to apply to the interconnection of content networks that are distributing streaming media.

2.4.2 Describing CDN Value: Reach and Scale

There are two fundamental elements that give a CDN value: outsourcing infrastructure and improved content delivery. A CDN allows multiple surrogates to act on behalf of an origin server, therefore removing the delivery of content from a centralized site to multiple and (usually) highly distributed sites. We refer to increased aggregate infrastructure size as "scale." In addition, a CDN can be constructed with copies of content near to end users, overcoming issues of network size, network congestion, and network failures. We refer to increased diversity of content locations as "reach."

Day, et al.	Expires August 23, 2002	[Page 9]
Internet-Draft	CDI Model	February 2002

In a typical (non-internetworked) CDN, a single service provider operates the request-routers, the surrogates, and the content distributors. In addition, that service provider establishes (business) relationships with content publishers and acts on behalf of their origin sites to provide a distributed delivery system. The value of that CDN to a content provider is a combination of its scale and its reach.

Day, et al.	Expires August 23, 2002	[Page 10]
Internet-Draft	CDI Model	February 2002

3. Content Network Model Terms

This section consists of the definitions of a number of terms used to refer to roles, participants, and objects involved in content networks. Although the following uses many terms that are based on those used in RFC 2616 [1] or RFC 3040 [3], there is no necessary connection to HTTP or web caching technology. Content internetworking and this vocabulary are applicable to other protocols and styles of content delivery.

Phrases in upper-case refer to other defined terms.

ACCOUNTING

Measurement and recording of DISTRIBUTION and DELIVERY activities, especially when the information recorded is ultimately used as a basis for the subsequent transfer of money, goods, or obligations.

ACCOUNTING SYSTEM

A collection of CONTENT NETWORK ELEMENTS that supports ACCOUNTING for a single CONTENT NETWORK.

AUTHORITATIVE REQUEST-ROUTING SYSTEM

The REQUEST-ROUTING SYSTEM that is the correct/final authority for a particular item of CONTENT.

CDN

Content Delivery Network or Content Distribution Network. A type of CONTENT NETWORK in which the CONTENT NETWORK ELEMENTS are arranged for more effective delivery of CONTENT to CLIENTS. Typically a CDN consists of a REQUEST-ROUTING SYSTEM, SURROGATES, a DISTRIBUTION SYSTEM, and an ACCOUNTING SYSTEM.

CLIENT

A program that sends CONTENT REQUESTS and receives corresponding CONTENT RESPONSES. [Note: this is similar to the definition in RFC 2616 [1] but we do not require establishment of a connection.]

CONTENT

Any form of digital data, CONTENT approximately corresponds to what is referred to as an "entity" in RFC 2616 [1]. One important form of CONTENT with additional constraints on DISTRIBUTION and DELIVERY is CONTINUOUS MEDIA.

CONTENT NETWORK

An arrangement of CONTENT NETWORK ELEMENTS, controlled by a common management in some fashion.

CONTENT NETWORK ELEMENT

Day, et al.	Expires August 23, 2002	[Page 11]
Internet-Draft	CDI Model	February 2002

A network device that performs at least some of its processing by examining CONTENT-related parts of network messages. In IP-based networks, a CONTENT NETWORK ELEMENT is a device whose processing depends on examining information contained in IP packet bodies; network elements (as defined in RFC 3040) examine only the header of an IP packet. Note that many CONTENT NETWORK ELEMENTS do not examine or even see individual IP packets, instead receiving the body of one or more packets assembled into a message of some higher-level protocol.

CONTENT REQUEST

A message identifying a particular item of CONTENT to be delivered.

CONTENT RESPONSE

A message containing a particular item of CONTENT, identified in a previous CONTENT REQUEST.

CONTENT SIGNAL

A message delivered through a DISTRIBUTION SYSTEM that specifies information about an item of CONTENT. For example, a CONTENT SIGNAL can indicate that the ORIGIN has a new version of some piece of CONTENT.

CONTINUOUS MEDIA

CONTENT where there is a timing relationship between source and sink; that is, the sink must reproduce the timing relationship that existed at the source. The most common examples of CONTINUOUS MEDIA are audio and motion video. CONTINUOUS MEDIA can be real-time (interactive), where there is a "tight" timing relationship between source and sink, or streaming (playback), where the relationship is less strict. [Note: This definition is essentially identical to the definition of continuous media in [2]]

DELIVERY

The activity of providing a PUBLISHER's CONTENT, via CONTENT RESPONSES, to a CLIENT. Contrast with DISTRIBUTION and REQUEST-ROUTING.

DISTRIBUTION

The activity of moving a PUBLISHER's CONTENT from its ORIGIN to one or more SURROGATES. DISTRIBUTION can happen either in anticipation of a SURROGATE receiving a REQUEST (pre-positioning) or in response to a SURROGATE receiving a REQUEST (fetching on demand). Contrast with DELIVERY and REQUEST-ROUTING.

DISTRIBUTION SYSTEM

Day, et al.	Expires August 23, 2002	[Page 12]
Internet-Draft	CDI Model	February 2002

A collection of CONTENT NETWORK ELEMENTS that support DISTRIBUTION for a single CONTENT NETWORK. The DISTRIBUTION SYSTEM also propagates CONTENT SIGNALS.

ORIGIN

The point at which CONTENT first enters a DISTRIBUTION SYSTEM.
The ORIGIN for any item of CONTENT is the server or set of servers at the "core" of the distribution, holding the "master" or "authoritative" copy of that CONTENT. [Note: We believe this definition is compatible with that for "origin server" in RFC 2616 [1] but includes additional constraints useful for CDI.]

PUBLISHER

The party that ultimately controls the CONTENT and its distribution.

REACHABLE SURROGATES

The collection of SURROGATES that can be contacted via a particular DISTRIBUTION SYSTEM or REQUEST-ROUTING SYSTEM.

REQUEST-ROUTING

The activity of steering or directing a CONTENT REQUEST from a USER AGENT to a suitable SURROGATE.

REQUEST-ROUTING SYSTEM

A collection of CONTENT NETWORK ELEMENTS that support REQUEST-ROUTING for a single CONTENT NETWORK.

SERVER

A program that accepts CONTENT REQUESTS and services them by sending back CONTENT RESPONSES. Any given program may be capable of being both a client and a server; our use of these terms refers only to the role being performed by the program. [Note: this is adapted from a similar definition in RFC 2616 [1].]

SURROGATE

A delivery server, other than the ORIGIN. Receives a CONTENT REQUEST and delivers the corresponding CONTENT RESPONSE. [Note: this is a different definition from that in RFC 3040 [3], which appears overly elaborate for our purposes. A "CDI surrogate" is always an "RFC 3040 surrogate"; we are not sure if the reverse is true.]

USER AGENT

The CLIENT which initiates a REQUEST. These are often browsers, editors, spiders (web-traversing robots), or other end user tools. [Note: this definition is identical to the one in RFC 2616 [1].]

Day, et al.	Expires August 23, 2002	[Page 13]
Internet-Draft	CDI Model	February 2002

4. Content Internetworking

There are limits to how large any one network's scale and reach can be. Increasing either scale or reach is ultimately limited by the cost of equipment, the space available for deploying equipment, and/or the demand for that scale/reach of infrastructure. Sometimes a particular audience is tied to a single service provider or a small set of providers by constraints of technology, economics, or law. Other times, a network provider may be able to manage surrogates and a distribution system, but may have no direct relationship with content providers. Such a provider wants to have a means of affiliating their delivery and distribution infrastructure with other parties who have content to distribute.

Content internetworking allows different content networks to share resources so as to provide larger scale and/or reach to each participant than they could otherwise achieve. By using commonly defined protocols for content internetworking, each content network can treat neighboring content networks as "black boxes", allowing them to hide internal details from each other.

5. Content Internetworking Model Terms

This section consists of the definitions of a number of terms used to refer to roles, participants, and objects involved in internetworking content networks. The purpose of this section is to identify common terms and provide short definitions. A more detailed technical discussion of these terms and their relationships appears in "Content Internetworking Architectural Overview" [4].

ACCOUNTING INTERNETWORKING

Interconnection of two or more ACCOUNTING SYSTEMS so as to enable the exchange of information between them. The form of ACCOUNTING INTERNETWORKING required may depend on the nature of the NEGOTIATED RELATIONSHIP between the peering parties -- in particular, on the value of the economic exchanges anticipated.

ADVERTISEMENT

Information about resources available to other CONTENT NETWORKS, exchanged via CONTENT INTERNETWORKING GATEWAYS. Types of ADVERTISEMENT include AREA ADVERTISEMENTS, CONTENT ADVERTISEMENTS, and DISTRIBUTION ADVERTISEMENTS.

AREA ADVERTISEMENT

ADVERTISEMENT from a CONTENT NETWORK's REQUEST-ROUTING SYSTEM about aspects of topology, geography and performance of a CONTENT NETWORK. Contrast with CONTENT ADVERTISEMENT, DISTRIBUTION ADVERTISEMENT.

BILLING ORGANIZATION

An entity that operates an ACCOUNTING SYSTEM to support billing within a NEGOTIATED RELATIONSHIP with a PUBLISHER.

CONTENT ADVERTISEMENT

ADVERTISEMENT from a CONTENT NETWORK's REQUEST-ROUTING SYSTEM about the availability of one or more collections of CONTENT on a CONTENT NETWORK. Contrast with AREA ADVERTISEMENT, DISTRIBUTION ADVERTISEMENT

CONTENT DESTINATION

A CONTENT NETWORK or DISTRIBUTION SYSTEM that is accepting CONTENT from another such network or system. Contrast with CONTENT SOURCE.

CONTENT INTERNETWORKING GATEWAY (CIG)

An identifiable element or system through which a CONTENT NETWORK can be interconnected with others. A CIG may be the point of contact for DISTRIBUTION INTERNETWORKING, REQUEST-ROUTING INTERNETWORKING, and/or ACCOUNTING INTERNETWORKING, and thus may

incorporate some or all of the corresponding systems for the CONTENT NETWORK.

CONTENT REPLICATION

The movement of CONTENT from a CONTENT SOURCE to a CONTENT DESTINATION. Note that this is specifically the movement of CONTENT from one network to another. There may be similar or different mechanisms that move CONTENT around within a single network's DISTRIBUTION SYSTEM.

CONTENT SOURCE

A CONTENT NETWORK or DISTRIBUTION SYSTEM that is distributing CONTENT to another such network or system. Contrast with CONTENT DESTINATION.

DISTRIBUTION ADVERTISEMENT

An ADVERTISEMENT from a CONTENT NETWORK's DISTRIBUTION SYSTEM to potential CONTENT SOURCES, describing the capabilities of one or more CONTENT DESTINATIONS. Contrast with AREA ADVERTISEMENT, CONTENT ADVERTISEMENT.

DISTRIBUTION INTERNETWORKING

Interconnection of two or more DISTRIBUTION SYSTEMS so as to propagate CONTENT SIGNALS and copies of CONTENT to groups of SURROGATES.

INJECTION

A "send-only" form of DISTRIBUTION INTERNETWORKING that takes place from an ORIGIN to a CONTENT DESTINATION.

INTER-

Describes activity that involves more than one CONTENT NETWORK (e.g. INTER-CDN). Contrast with INTRA-.

INTRA-

Describes activity within a single CONTENT NETWORK (e.g. INTRA-CDN). Contrast with INTER-.

NEGOTIATED RELATIONSHIP

A relationship whose terms and conditions are partially or completely established outside the context of CONTENT NETWORK internetworking protocols.

REMOTE CONTENT NETWORK

A CONTENT NETWORK able to deliver CONTENT for a particular REQUEST that is not the AUTHORITATIVE REQUEST-ROUTING SYSTEM for that REQUEST.

Day, et al.	Expires August 23, 2002	[Page 16]
Internet-Draft	CDI Model	February 2002

REQUEST-ROUTING INTERNETWORKING

Interconnection of two or more REQUEST-ROUTING SYSTEMS so as to increase the number of REACHABLE SURROGATES for at least one of the interconnected systems.

Day, et al.	Expires August 23, 2002	[Page 17]
Internet-Draft	CDI Model	February 2002

6. Security Considerations

There are no security-related issues related to the terms defined in this document. The technology of content internetworking does raise some security-related issues, and a detailed discussion of those issues appears in "Content Internetworking Architectural Overview" [4].

7. Acknowledgements

The authors acknowledge the contributions and comments of Fred Douglass (AT&T), Don Gilletti (CacheFlow), Markus Hoffmann (Lucent), Barron Housel (Cisco), Barbara Liskov (Cisco), John Martin (Network Appliance), Nalin Mistry (Nortel Networks) Raj Nair (Cisco), Hilarie Orman (Volera), Doug Potter (Cisco), and Oliver Spatscheck (AT&T).

[Note to RFC Editor: The last normative reference is [4], all subsequent references starting with [5] can be deleted.]

References

- [1] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P. and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2616, June 1999, <<http://www.rfc-editor.org/rfc/rfc2616.txt>>.
- [2] Schulzrinne, H., Rao, A. and R. Lanphier, "Real Time Streaming Protocol", RFC 2326, April 1998, <<http://www.rfc-editor.org/rfc/rfc2326.txt>>.
- [3] Cooper, I., Melve, I. and G. Tomlinson, "Internet Web Replication and Caching Taxonomy", RFC 3040, June 2000, <<http://www.rfc-editor.org/rfc/rfc3040.txt>>.
- [4] Green, M., Cain, B., Tomlinson, G., Thomas, S. and P. Rzewskip, "Content Internetworking Architectural Overview", draft-ietf-cdi-architecture-00.txt (work in progress), February 2002, <<http://www.ietf.org/internet-drafts/draft-ietf-cdi-architecture-00.txt>>.
- [5] Day, M., Gilletti, D. and P. Rzewski, "Content Internetworking Scenarios", draft-ietf-cdi-scenarios-00.txt (work in progress), February 2002, <<http://www.ietf.org/internet-drafts/draft-ietf-cdi-scenarios-00.txt>>.
- [6] Gilletti, D., Nair, R., Scharber, J. and J. Guha, "Content Internetworking (CDI) Authentication, Authorization, and Accounting Requirements", draft-ietf-cdi-aaa-reqs-00.txt (work in progress), February 2002, <<http://www.ietf.org/internet-drafts/draft-ietf-cdi-aaa-reqs-00.txt>>.
- [7] Barbir, A., Cain, B., Douglass, F., Green, M., Hoffmann, M., Nair, R., Potter, D. and O. Spatscheck, "Known CDN Request-Routing Mechanisms", draft-ietf-cdi-known-request-routing-00.txt (work in progress), February 2002, <<http://www.ietf.org/internet-drafts/draft-ietf-cdi-known-request-routing-00.txt>>.
- [8] Cain, B., Spatscheck, O., May, M. and A. Barbir, "Request-Routing Requirements for Content Internetworking", draft-ietf-cdi-request-routing-reqs-00.txt (work in progress), February 2002, <<http://www.ietf.org/internet-drafts/draft-ietf-cdi-request-routing-reqs-00.txt>>.
- [9] Amini, L., Spatscheck, O. and S. Thomas, "Distribution Requirements for Content Internetworking", draft-ietf-cdi-distribution-reqs-00.txt (work in progress), February 2002, <<http://www.ietf.org/internet-drafts/draft-ietf-cdi-distribution-reqs-00.txt>>.

Day, et al. Expires August 23, 2002 [Page 20]

Internet-Draft CDI Model February 2002

distribution-reqs-00.txt>.

Authors' Addresses

Mark Stuart Day
Cisco Systems
1414 Massachusetts Avenue
Boxborough, MA 01719
US

Phone: +1 978 936 1089
EMail: markday@cisco.com

Brad Cain
Cereva Networks
3 Network Drive
Marlborough, MA 01752
US

Phone: +1 508-787-5000
EMail: bcain@cereva.com

Gary Tomlinson
CacheFlow, Inc.
12034 134th Ct. NE Suite 201
Redmond, WA 98052
US

Phone: +1 425 820 3009
EMail: garyt@cacheflow.com

Phil Rzewski
Inktomi
4100 East Third Avenue, MS FC2-4
Foster City, CA 94404
US

Day, et al. Expires August 23, 2002 [Page 21]
Internet-Draft CDI Model February 2002

Full Copyright Statement

Copyright (C) The Internet Society (2002). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

Day, et al. Expires August 23, 2002 [Page 22]

EXHIBIT 2

Content Internetworking Architectural Overview
draft-green-cdn-p-gen-arch-03.txt

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of RFC2026.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on August 31, 2001.

Copyright Notice

Copyright (C) The Internet Society (2001). All Rights Reserved.

Abstract

There is wide interest in the technology for interconnecting Content Networks, variously called "Content Peering" or "Content Internetworking". We present the general architecture and core building blocks used in the internetworking of Content Networks. The scope of this work is limited to external interconnections with Content Networks and does not address internal mechanisms used within Content Networks, which for the purpose of the document are considered to be black boxes. This work establishes an abstract architectural framework to be used in the development of protocols, interfaces, and system models for standardized Content Internetworking.

Table of Contents

1.	Introduction	4
2.	Content Internetworking System Architecture	5
2.1	Conceptual View of Peered Content Networks	5
2.2	Content Internetworking Architectural Elements	7
3.	Request-Routing Peering System	11
3.1	Request-Routing Overview	11
3.2	Request Routing	13
3.3	System Requirements	13
3.4	Protocol Requirements	14
3.5	Examples	14
3.6	Request-Routing Problems to Solve	15
4.	Distribution Peering System	17
4.1	Distribution Overview	17
4.2	Distribution Models	19
4.3	Distribution Components	20
4.4	Distribution System Requirements	20
4.4.1	Replication Requirements	21
4.4.2	Signaling Requirements	21
4.4.3	Advertising Requirements	21
4.5	Protocol Requirements	22
4.6	Distribution Problems to Solve	22

4.6.1	General Problems	22
4.6.2	Replication Problems	23
4.6.3	Signaling Problems	23
4.6.4	Advertising Problems	23
5.	Accounting Peering System	25
5.1	Accounting Overview	25
5.2	Accounting System Requirements	27
5.3	Protocol Requirements	27
6.	Security Considerations	28
6.1	Threats to Content Internetworking	28
6.1.1	Threats to the CLIENT	28
6.1.1.1	Defeat of CLIENT's Security Settings	28

6.1.1.2	Delivery of Bad Accounting Information	28
6.1.1.3	Delivery of Bad CONTENT	29
6.1.1.4	Denial of Service	29
6.1.1.5	Exposure of Private Information	29
6.1.1.6	Substitution of Security Parameters	29
6.1.1.7	Substitution of Security Policies	29
6.1.2	Threats to the PUBLISHER	29
6.1.2.1	Delivery of Bad Accounting Information	29
6.1.2.2	Denial of Service	30
6.1.2.3	Substitution of Security Parameters	30
6.1.2.4	Substitution of Security Policies	30
6.1.3	Threats to a CN	30
6.1.3.1	Bad Accounting Information	30
6.1.3.2	Denial of Service	30
6.1.3.3	Transitive Threats	31
7.	Acknowledgements	32
	References	33
	Authors' Addresses	35
	Full Copyright Statement	36

1. Introduction

Terms in ALL CAPS, except those qualified with explicit citations are defined in [13].

This memo describes the overall architectural structure and the fundamental building blocks used in the composition of Content Internetworking. Consult [13] for the system model, and vocabulary used in, this application domain. A key requirement of the architecture itself is that it be able to address each of the Content Internetworking scenarios enumerated in [14]. The scope of this work is limited to external interconnections between Content Networks (CN) (i.e. INTER-CN) and does not address internal mechanisms used within Content Networks (i.e. INTRA-CN), which for the purposes of the document are considered to be black boxes. This work is intended to establish an abstract architectural framework to be used in the development of protocols, interfaces and system models for standardized, interoperable peering among Content Networks.

We first present the architecture as an abstract system. Then we develop a more concrete system architecture. For each core

architectural element, we first present the structure of the element followed by system requirements. Protocol requirements for individual core elements are presented in accompanying works [17][18][15]. The assumptions and scenarios constraining the architecture is explained in [14]. We intend that the architecture should support a wide variety of business models.

At the core of Content Internetworking are three principal architectural elements that constitute the building blocks of the Content Internetworking system. These elements are the REQUEST-ROUTING PEERING SYSTEM, DISTRIBUTION PEERING SYSTEM, and ACCOUNTING PEERING SYSTEM. Collectively, they control selection of the delivery Content Network, content distribution between peering Content Networks, and usage accounting, including billing settlement among the peering Content Networks.


This work takes into consideration relevant IETF RFCs and IETF works-in-progress. In particular, it is mindful of the end-to-end nature [6][10] of the Internet, the current taxonomy of web replication and caching [11], and the accounting, authorization and authentication framework [12].

2. Content Internetworking System Architecture

2.1 Conceptual View of Peered Content Networks

Before developing the system architecture, a conceptual view of peered CNs is presented to frame the problem space. CNs are comprised principally of four core system elements [13], the REQUEST-ROUTING SYSTEM, the DISTRIBUTION SYSTEM, the ACCOUNTING SYSTEM, and SURROGATES. In order for CNs to peer with one another, it is necessary to interconnect several of the core system elements of individual CNs. The interconnection of CN core system elements occurs through network elements called Content Peering Gateways (CPG). Namely, the CN core system elements that need to be interconnected are the REQUEST-ROUTING SYSTEM, the DISTRIBUTION SYSTEM, and the ACCOUNTING SYSTEM.

Figure 1 contains a conceptual peered Content Networks diagram.






Figure 1 Conceptual View of Peered Content Networks

This conceptual view illustrates the peering of three Content Networks; CN A, CN B, and CN C. The CNs are peered through interconnection at Content Peering Gateways. The result is presented as a virtual CN to CLIENTs for the DELIVERY of CONTENT by the aggregated set of SURROGATES.

Note:
Not all Content Networks contain the complete set of core elements. For these Content Networks, peering will be done with only the core elements they do contain.

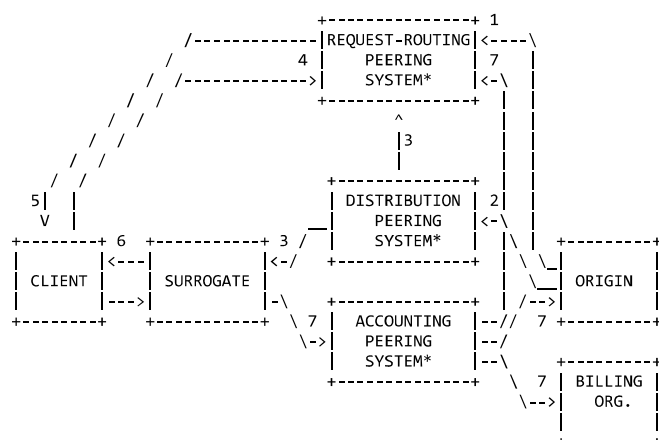
2.2 Content Internetworking Architectural Elements

The system architecture revolves around the general premise that individual Content Networks are wholly contained within an administrative domain [3] that is composed of either autonomous systems [1] (physical networks) or overlay networks (virtual networks). For the purpose of this memo, an overlay network is defined as a set of connected CN network elements layered onto existing underlying networks, and present the result as a virtual application layer to both CLIENTs and ORIGINS. The system architecture for CN peering accommodates this premise by assuring that the information and controls are available for inter-CN-domain administration. Content Internetworking involves the interconnection of the individual CN administrative domains through gateway protocols and mechanisms loosely modeled after BGP [5].

The system architecture depends on the following assumptions:

1. The URI [8] name space is the basis of PUBLISHER object identifiers.
2. PUBLISHERs delegate authority of their object URI name space being distributed by peering CNs to the REQUEST-ROUTING PEERING SYSTEM.
3. Peering CNs use a common convention for encoding CN metadata into the URI name space.

Figure 2 contains a system architecture diagram of the core elements involved in Content Internetworking.



Note: * represents core elements of Content Internetworking

Figure 2 System Architecture Elements of a Content Internetworking System

The System Architecture is comprised of 7 major elements, 3 of which constitute the Content Internetworking system itself. The peering elements are REQUEST-ROUTING PEERING SYSTEM, DISTRIBUTION PEERING SYSTEM, and ACCOUNTING PEERING SYSTEM. Correspondingly, the system architecture is a system of systems:

1. The ORIGIN delegates its URI name space for objects to be distributed and delivered by the peering CNS to the REQUEST-ROUTING PEERING SYSTEM.
2. The ORIGIN INJECTS CONTENT that is to be distributed and delivered by the peering CNS into the DISTRIBUTION PEERING SYSTEM.

Note:

CONTENT which is to be pre-populated (pushed) within the peering CNS is pro-actively injected, while CONTENT which is to be pulled on demand is injected at the time the object is being requested for DELIVERY.

3. The DISTRIBUTION PEERING SYSTEM moves content between CN DISTRIBUTION SYSTEMS. Additionally this system interacts with the REQUEST-ROUTING PEERING SYSTEM via feedback ADVERTISEMENTS to assist in the peered CN selection process for CLIENT requests.
4. The CLIENT requests CONTENT from what it perceives to be the ORIGIN, however due to URI name space delegation, the request is actually made to the REQUEST-ROUTING PEERING SYSTEM.

Note:

The request routing function may be implied by an in-path network element such as caching proxy, which is typical for a Access Content Network. In this case, request routing is optimized to a null function, since the CLIENT is a priori mapped to the SURROGATE.

5. The REQUEST-ROUTING PEERING SYSTEM routes the request to a suitable SURROGATE in a peering CN. REQUEST-ROUTING PEERING SYSTEMS interact with one another via feedback ADVERTISEMENTS in order to keep request-routing tables current.
6. The selected SURROGATE delivers the requested content to the CLIENT. Additionally, the SURROGATE sends accounting information for delivered content to the ACCOUNTING PEERING SYSTEM.
7. The ACCOUNTING PEERING SYSTEM aggregates and distills the accounting information into statistics and content detail records for use by the ORIGIN and BILLING ORGANIZATION. Statistics are also used as feedback to the REQUEST-ROUTING PEERING SYSTEM.

8. The BILLING ORGANIZATION uses the content detail records to settle with each of the parties involved in the content distribution and delivery process.

This process has been described in its simplest form in order to present the Content Internetworking architecture in the most abstract way possible. In practice, this process is more complex when applied to policies, business models and service level agreements that span multiple peering Content Networks. The orthogonal core peering systems are discussed in greater depth in Section 3, Section 4 and Section 5 respectively.

Green, et. al.	Expires August 31, 2001	[Page 9]
Internet-Draft	CDI Architecture	March 2001

Note:

Figure 2 simplifies the presentation of the core Content Internetworking elements as single boxes, when in fact they represent a collection of CPGs and interconnected individual CN core system elements. This has been done to introduce the system architecture at its meta level.

The system architecture does not impose any administrative domain [3] restrictions on the core peering elements (REQUEST-ROUTING PEERING SYSTEM, DISTRIBUTION PEERING SYSTEM and ACCOUNTING PEERING SYSTEM). The only requirement is that they be authorized by the principal parties (ORIGIN and peering CNs) to act in their behalf. Thus, it is possible for each of the core elements to be provided by a different organization.

Green, et. al.	Expires August 31, 2001	[Page 10]
Internet-Draft	CDI Architecture	March 2001

3. Request-Routing Peering System

The REQUEST-ROUTING PEERING SYSTEM represents the request-routing function of the Content Internetworking system. It is responsible for routing CLIENT requests to an appropriate peered CN for the delivery of content.

Note:

When the DISTRIBUTION PEERING SYSTEM and/or the ACCOUNTING PEERING SYSTEM is present, it is highly desirable to utilize content location information within the peered CNs and/or system load information in the selection of appropriate peered CNs in the routing of requests.

3.1 Request-Routing Overview

REQUEST-ROUTING SYSTEMs route CLIENT requests to a suitable SURROGATE, which is able to service a client request. Many

request-routing systems route users to the surrogate that is "closest" to the requesting user, or to the "least loaded" surrogate. However, the only requirement of the request-routing system is that it route users to a surrogate that can serve the requested content.

REQUEST-ROUTING PEERING is the interconnection of two or more REQUEST-ROUTING SYSTEMs so as to increase the number of REACHABLE SURROGATES for at least one of the interconnected systems.

In order for a PUBLISHER's CONTENT to be delivered by multiple peering CNS, it is necessary to federate each Content Network REQUEST-ROUTING SYSTEM under the URI name space of the PUBLISHER object. This federation is accomplished by first delegating authority of the PUBLISHER URI name space to an AUTHORITATIVE REQUEST-ROUTING SYSTEM. The AUTHORITATIVE REQUEST-ROUTING SYSTEM subsequently splices each peering Content Network REQUEST-ROUTING SYSTEM into this URI name space and transitively delegates URI name space authority to them for their participation in request-routing. Figure 3 is a diagram of the entities involved in the REQUEST-ROUTING PEERING SYSTEM.

Note:
 For the null request routing case (in path caching proxy present), the caching proxy acts as the SURROGATE. In this case, the SURROGATE performs the request routing via its pre-established proxy relationship with the CLIENT and is implicitly the terminating level of request routing. In essence, the SURROGATE is federated into the URI namespace without the need to communicate with the AUTHORITATIVE REQUEST-ROUTING SYSTEM.

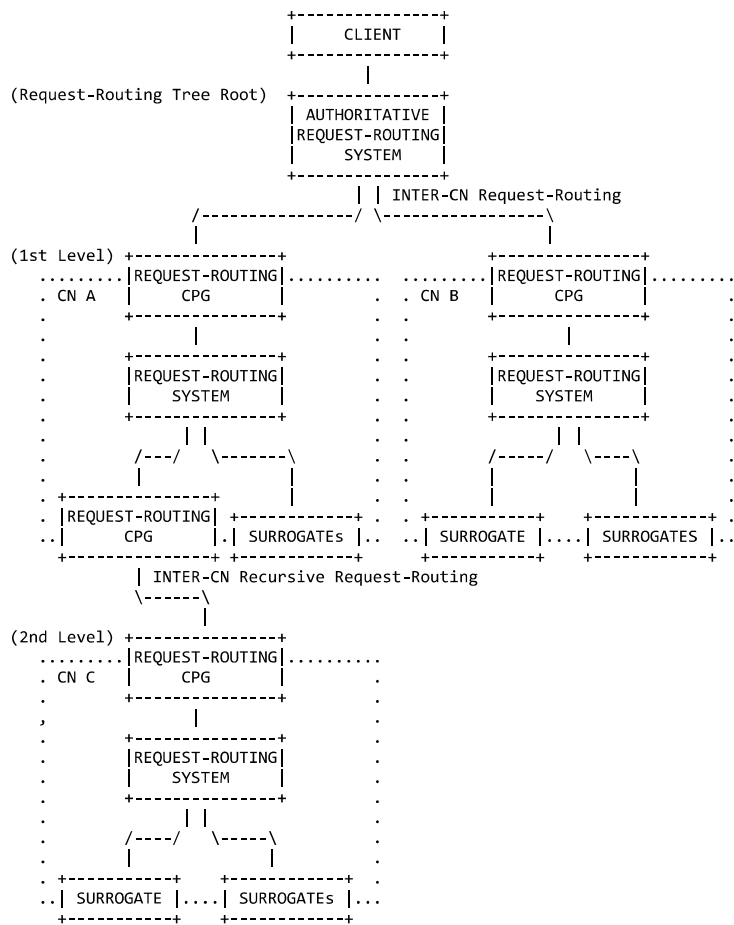


Figure 3 REQUEST-ROUTING PEERING SYSTEM Architecture

The REQUEST-ROUTING PEERING SYSTEM is hierarchical in nature. There exists exactly one request-routing tree for each PUBLISHER URI. The AUTHORITATIVE REQUEST-ROUTING SYSTEM is the root of the

request-routing tree. There may be only one AUTHORITATIVE REQUEST-ROUTING SYSTEM for a URI request-routing tree. Subordinate to the AUTHORITATIVE REQUEST-ROUTING SYSTEM are the REQUEST-ROUTING SYSTEMS of the first level peering CNs. There may exist recursive subordinate REQUEST-ROUTING SYSTEMS of additional level peering CNs.

Note:

A PUBLISHER object may have more than one URI associated with it and therefore be present in more than one request-routing tree.

3.2 Request Routing

The actual "routing" of a client request is through REQUEST-ROUTING CPGs. The AUTHORITATIVE REQUEST-ROUTING CPG receives the CLIENT request and forwards the REQUEST to an appropriate DISTRIBUTING CN. This process of INTER-CN request-routing may occur multiple times in a recursive manner between REQUEST-ROUTING CPGs until the REQUEST-ROUTING SYSTEM arrives at an appropriate DISTRIBUTING CN to deliver the content.

Note:

The Client request may be for resolution of a URI component and not the content of the URI itself. This is the case when DNS is being utilized in the request-routing process to resolve the URI server component.

Request-Routing systems explicitly peer but do not have "interior" knowledge of surrogates from other CNs. Each CN operates its internal request-routing system. In this manner, request-routing systems peer very much like IP network layer peering.

3.3 System Requirements

We assume that there is a peering relationship between REQUEST-ROUTING CPGs. This peering relationship at a minimum must exchange a set of CLIENT IP addresses that can be serviced, and a set of information about the DISTRIBUTION SYSTEMS, for which they are performing request-routing.

Green, et. al.	Expires August 31, 2001	[Page 13]
Internet-Draft	CDI Architecture	March 2001

Request-Routing Requirements

1. Use of a URI name space based request-routing mechanism. The request-routing mechanism is allowed to use as much of the URI name space as it needs to select the proper SURROGATE. For example, DNS based mechanisms utilize only the host subcomponent, while content aware mechanisms utilize use multiple components.
2. Normalized canonical URI name space structure for peered CN distribution of PUBLISHER objects. The default in the absence of encoded meta data is the standard components as defined by [8]. Encoded meta data must conform to the syntactical grammar defined in [7].
3. Single AUTHORITATIVE REQUEST-ROUTING SYSTEM for PUBLISHER object URI name space.
4. Assure that the request-routing tree remains a tree -- i.e., has no cycles.
5. Assure that adjacent request-routing systems from different administrative domains (different CNs) use a compatible request-routing mechanism.
6. Assure that adjacent request-routing systems from different administrative domains (different CNs) agree to forward requests for the CONTENT in question.
- 7.

[Editor Note:

System requirements being generated in the request-routing peering protocol design team have not yet been reconciled and integrated into this document.]

3.4 Protocol Requirements

Consult [17] for request-routing peering protocol requirements.

3.5 Examples

Consult [16] for in-depth information on known request-routing systems.

3.6 Request-Routing Problems to Solve

[Editor Note:

This section is being preserved until it has been determined that these issues have been addressed in the request-routing peering protocol requirements draft.]

Specific problems in request-routing needing further investigation include:

1. What is the aggregated granularity of CLIENT IP address being serviced by a peering CN's DISTRIBUTION SYSTEM?
2. How do DNS request-routing systems forward a request? If a given CN is peered with many other CNs, what are the criteria that forwards a request to another CN?
3. How do content-aware request-routing systems forward a request? If a given CN is peered with many other CNs, what are the criteria that forwards a request to another CN?
4. What are the merits of designing a generalized content routing protocol, rather than relying on request-routing mechanisms.
5. What is the normalized canonical URI name space for request-routing? Because request-routing is federated across multiple CNs, it is necessary to have agreed upon standards for the encoding of meta data in URIs. There are many potential elements, which may be encoded. Some of these elements are: authoritative agent domain, publisher domain, content type, content length, etc.
6. How are policies communicated between the REQUEST-ROUTING SYSTEM and the DISTRIBUTION ADVERTISEMENT SYSTEM? A given CN may wish to serve only a given content type or a particular set of users. These types of policies must be communicated between CNs.
7. What are the request-routing protocols in DNS? When a request is routed to a particular REQUEST-ROUTING CPG, a clear set of DNS rules and policies must be followed in order to have a workable and predictable system.
8. How do we protect the REQUEST-ROUTING SYSTEM against denial of service attacks?

9. How do we select the appropriate peering CN for DELIVERY?

The selection process must to consider the distribution policies involved in Section 4. Investigation into other policy "work in progress" within the IETF is needed to understand the relationship of policies developed within Content Internetworking.

4. Distribution Peering System

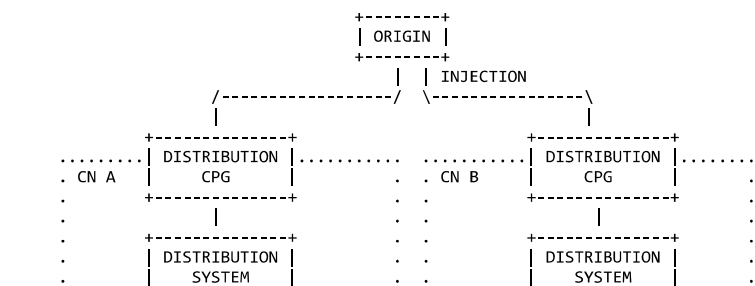
The DISTRIBUTION PEERING SYSTEM represents the content distribution function of the CN peering system. It is responsible for moving content from one DISTRIBUTION CPG to another DISTRIBUTION CPG and for supplying content location information to the REQUEST-ROUTING PEERING SYSTEM.

4.1 Distribution Overview

One goal of the Content Internetworking system is to move content closer to the CLIENT. Typically this is accomplished by copying content from its ORIGIN to SURROGATES. The SURROGATES then have the CONTENT available when it is requested by a CLIENT. Even with a single PUBLISHER and single CN, the copying of CONTENT to a SURROGATE may traverse a number of links, some in the PUBLISHER's network, some in the CN's network, and some between those two networks. For DISTRIBUTION PEERING, we consider only the communication "between" two networks, and ignore the mechanisms for copying CONTENT within a network.

In the above example the last server on the content provider's network in the path, and the first server on the CN's network in the path, must contain DISTRIBUTION CPGs which communicate directly with each other. The DISTRIBUTION CPGs could be located in the ORIGIN server and the SURROGATE server. Thus in the simplest form the ORIGIN server is in direct contact with the SURROGATE. However the DISTRIBUTION CPG in the content provider's network could aggregate content from multiple ORIGIN servers and the DISTRIBUTION CPG in the CN's network could represent multiple SURROGATES. These DISTRIBUTION CPGs could then be co-located in an exchange facility. In fact, given the common practice of independently managed IP peering co-location exchange facilities for layer 3, there exists the distinct opportunity to create similar exchanges for CPGs.

Figure 4 is a diagram of the entities involved in the DISTRIBUTION PEERING SYSTEM.



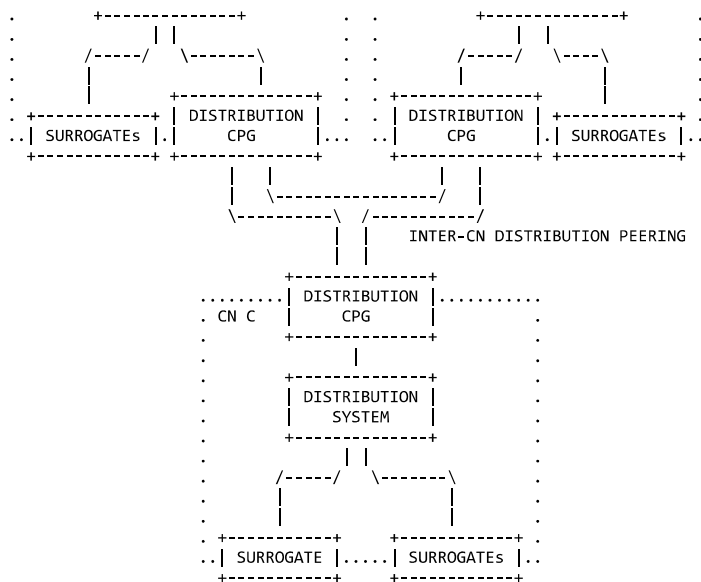


Figure 4 DISTRIBUTION PEERING SYSTEM Architecture

While Content Internetworking in general relates to interfacing with CNs, there are two CN distribution peering relationships we expect to be common; INTER-CN distribution peering and INJECTION peering. INTER-CN distribution peering involves distributing CONTENT between individual CNs in a inter-network of peered CNs. INJECTION peering involves the publishing of CONTENT directly into CNs by ORIGINs.

4.2 Distribution Models

Replication ADVERTISEMENTS may take place in a model similar to the way IP routing table updates are done between BGP routers. DISTRIBUTION CPGs could take care of exterior content replication between content providers and CNs, while at the same time performing content replication interior to their networks in an independent manner. If this model is used then the internal structure of the networks is hidden and the only knowledge of other networks is the locations of DISTRIBUTION CPGs.

Replication of content may take place using a push model, or a pull model, or a combination of both. Use initiated replication, where SURROGATES, upon getting a cache miss, retrieve CONTENT from the DISTRIBUTION SYSTEM, represents the pull model. ORIGIN initiated replication of CONTENT to SURROGATES represents the push model. DISTRIBUTION CPGs may be located at various points in these models depending on the topologies of the networks involved.

With Content Internetworking it may be desirable to replicate content through a network, which has no internal SURROGATES. For example add a exchange network between the content provider network and the CN network to the example above. The exchange network could have a DISTRIBUTION CPG co-located with the content provider's DISTRIBUTION CPG, which acts as a proxy for the CN. The exchange network could also have a DISTRIBUTION CPG co-located with the CN's DISTRIBUTION CPG, which acts as a proxy for the content provider. In a consolidated example, the exchange network could have a single DISTRIBUTION CPG that acts as a proxy for both the content provider and the CN.

Replication of CONTINUOUS MEDIA that is not to be cached on SURROGATES, such as live streaming broadcasts, takes place in a different model from content that is to be persistently stored. Replication in this case, typically takes the form of splitting the live streaming data at various points in the network. In Content Internetworking, DISTRIBUTION CPGs may support CONTINUOUS MEDIA splitting replication, as they likely provide ideal network topologic points for application layer multicasting.

4.3 Distribution Components

The three main components of DISTRIBUTION PEERING are replication, signaling and advertising.

The first component of content distribution is replication. Replication involves moving the content from an ORIGIN server to SURROGATE servers. The immediate goal in CN peering is moving the content between DISTRIBUTION CPGs.

The second component of content distribution is content signaling. Content signaling is the propagation of content meta-data. This meta-data may include such information such as the immediate expiration of content or a change in the expiration time of CONTENT. The immediate goal in signaling is exchanging signals between DISTRIBUTION CPGs.

The third component of content distribution is content advertising. Content providers must be able to advertise content that can be distributed by CNs and its associated terms. It is important that the advertising of content must be able to aggregate content information. The immediate goal in advertising is exchanging advertisements between DISTRIBUTION CPGs.

4.4 Distribution System Requirements

Replication systems must have a peering relationship. This peering relationship must exchange sets of aggregated content and its meta-data. Meta-data may change over time independently of the content data and must be exchanged independently as well.

Green, et. al.	Expires August 31, 2001	[Page 20]
Internet-Draft	CDI Architecture	March 2001

4.4.1 Replication Requirements

The specific requirements in content replication are:

1. A common protocol for the replication of content.
2. A common format for the actual content data in the protocol.
3. A common format for the content meta-data in the protocol.
4. Security mechanisms (see Section 6).
5. Scalable distribution of the content.

4.4.2 Signaling Requirements

The specific requirements in content signaling are:

1. Signals for (at least) "flush" and "expiration time update".
2. Security mechanisms (see Section 6).
3. Scalable distribution of the signals on a large scale.

Editor Note:

We have to start being quantitative about what we mean by "large scale". Are we thinking in terms of the number of content items, the number of networks, or the number of signals? For each of those, how big is "large scale"?

4. Content location and serviced CLIENT IP aggregate address exchanges with REQUEST-ROUTING CPGs.

4.4.3 Advertising Requirements

The specific requirements in CONTENT ADVERTISEMENT are:

1. A common protocol for the ADVERTISEMENT of CONTENT.
2. A common format for the actual ADVERTISEMENTS in the protocol.

Editor Note:

The following requirements need further discussion. As it stands now, there isn't sufficient information to substantiate them.

3. A well-known state machine.
4. Use of TCP or SCTP (because soft-state protocols will not scale).

Green, et. al. Expires August 31, 2001 [Page 21]
Internet-Draft CDI Architecture March 2001

5. Well-known error codes to diagnose protocols between different networks.
6. Capability negotiation.
7. Ability to represent policy.

[Editor Note:

System requirements being generated in the distribution peering protocol design team have not yet been reconciled and integrated into this document.]

4.5 Protocol Requirements

Consult [18] for distribution peering protocol requirements.

4.6 Distribution Problems to Solve

[Editor Note:

This section is being preserved until it has been determined that these issues have been addressed in the distribution peering protocol requirements draft.]

Some of the problems in distribution revolve around supporting both a push model and a pull model for replication of content in that they are not symmetric. The push model is used for pre-loading of content and the pull model is used for on-demand fetching and pre-fetching of content. These models are not symmetric in that the amount of available resources in which to place the content on the target server must be known. In the fetching cases the server that pulls the content knows the available resources on the target server, itself. In the pre-loading case the server that pushes the content must find out the available resources from the target server before pushing the data.

4.6.1 General Problems

General problems in distribution peering needing further investigation include:

1. How would a single distribution peering protocol adequately support replication, signaling and advertising?
2. Should a single distribution peering protocol be considered, rather than separate protocols for each component?
3. How do we prevent looping of distribution updates? That is to say, detect and stop propagating replication, signaling and advertisement of events a DISTRIBUTION CPG has already issued.

Green, et. al. Expires August 31, 2001 [Page 22]
Internet-Draft CDI Architecture March 2001

Looping here has the possibility of becoming infinite, if not bounded by the protocol(s). IP route updating and forwarding has faced similar issues and has solved them.

4.6.2 Replication Problems

Specific problems in replication needing further investigation include:

1. How do replication systems forward a request?
2. How do we keep pull based replication serviced within the DISTRIBUTION CPGs in order to prevent it from inadvertently bleeding out into REQUEST-ROUTING SYSTEM and potentially getting into a recursive loop?
3. How are policies communicated between the replication systems?
4. What are the replication protocols?
5. Does replication only take place between CPGs?

4.6.3 Signaling Problems

Specific problems in content signaling needing further investigation

include:

1. How do we represent a content signal?
2. What content meta-data needs to be signaled?
3. How do we represent aggregates of meta-data in a concise and compressed manner?
4. What protocol(s) should be used for content signals?
5. What is a scalable architecture for delivering content signals?
6. Do content signals need a virtual distribution system of their own?

4.6.4 Advertising Problems

Specific problems in CONTENT ADVERTISEMENT needing further investigation include:

1. How do we represent aggregates of content to be distributed in a concise and compressed manner?

Green, et. al.	Expires August 31, 2001	[Page 23]
Internet-Draft	CDI Architecture	March 2001

2. What protocol(s) should be used for the aggregation of this data?
3. What are the issues involved in the creation of CPG exchanges?
This is actually a broader question than just for distribution, but needs to be considered for all forms of CPGs {REQUEST-ROUTING, DISTRIBUTION, ACCOUNTING}.

Green, et. al.	Expires August 31, 2001	[Page 24]
Internet-Draft	CDI Architecture	March 2001

5. Accounting Peering System

The ACCOUNTING PEERING SYSTEM represents the accounting data collection function of the Content Internetworking system. It is responsible for moving accounting data from one ACCOUNTING CPG to another ACCOUNTING CPG.

5.1 Accounting Overview

Content Internetworking must provide the ability for the content

provider to collect data regarding the delivery of their CONTENT by the peered CNS. ACCOUNTING CPGs exchange the data collected by the interior ACCOUNTING SYSTEMS. This interior data may be collected from the SURROGATES by ACCOUNTING CPGs using SNMP or FTP, for example. ACCOUNTING CPGs may transfer the data to exterior neighboring ACCOUNTING CPGs on request (push), in an asynchronous manner (push), or a combination of both. Accounting data may also be aggregated before it is transferred.

Figure 5 is a diagram of the entities involved in the ACCOUNTING PEERING SYSTEM.

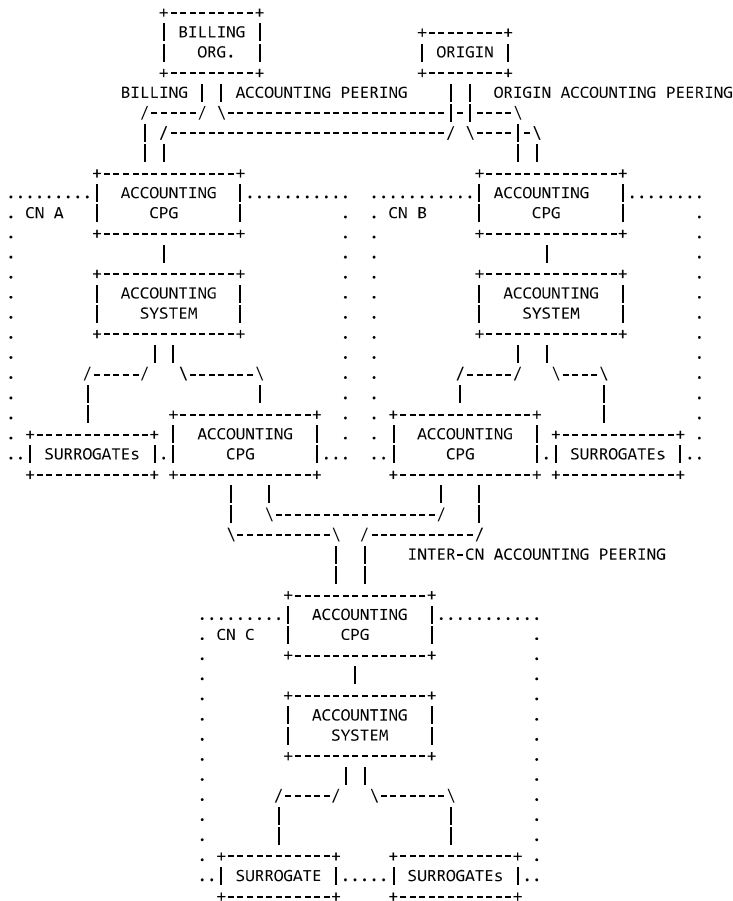


Figure 5 ACCOUNTING Peering system Architecture

There are three CN accounting peering relationships we expect to be common; INTER-CN accounting peering, BILLING ORGANIZATION accounting peering and ORIGIN accounting peering. INTER-CN accounting peering involves exchanging accounting information between individual CNs in a inter-network of peered CNs. BILLING ORGANIZATION peering involves exchanging to accounting information between CNs and a billing organization. ORIGIN accounting peering involves the exchanging of accounting information between CNs and ORIGINS.

Note:

It is not necessary for an ORIGIN to peer directly with multiple CNs in order to participate in Content Internetworking. ORIGINS participating in a single home CN will be indirectly peered by their home CN with the inter-network of CNs the home CN is a member of. Nor is it necessary to have a BILLING ORGANIZATION peer, since this function may also be provided by the home CN. However, ORIGINS that directly peer for ACCOUNTING may have access to greater accounting detail. Also, through the use of ACCOUNTING peering, 3rd party billing can be provided.

5.2 Accounting System Requirements

[Editor Note:

System requirements being generated in the accounting peering protocol design team have not yet been reconciled and integrated into this document.]

5.3 Protocol Requirements

Consult [15] for accounting peering protocol requirements.

6. Security Considerations

Security concerns with respect to Content Internetworking can be generally categorized into trust within the system and protection of the system from threats. The trust model utilized with Content Internetworking is predicated largely on transitive trust between the ORIGIN, REQUEST-ROUTING PEERING SYSTEM, DISTRIBUTION PEERING SYSTEM, ACCOUNTING PEERING SYSTEM and SURROGATES. Network elements within the Content Internetworking system are considered to be "insiders" and therefore trusted.

6.1 Threats to Content Internetworking

The following sections document key threats to CLIENTs, PUBLISHERs, and CNs. The threats are classified according to the party that they most directly harm, but, of course, a threat to any party is ultimately a threat to all. (For example, having a credit card number stolen may most directly affect a CLIENT; however, the resulting dissatisfaction and publicity will almost certainly cause some harm to the PUBLISHER and CN, even if the harm is only to those organizations' reputations.)

6.1.1 Threats to the CLIENT

6.1.1.1 Defeat of CLIENT's Security Settings

Because the SURROGATE's location may differ from that of the ORIGIN, the use of a SURROGATE may inadvertently or maliciously defeat any location-based security settings employed by the CLIENT. And since the SURROGATE's location is generally transparent to the CLIENT, the CLIENT may be unaware that its protections are no longer in force. For example, a CN may relocate CONTENT from a Internet Explorer user's "Internet Web Content Zone" to that user's "Local Intranet Web Content Zone." If the relocation is visible to the Internet Explorer browser but otherwise invisible to the user, the browser may be employing less stringent security protections than the user

is expecting for that CONTENT. (Note that this threat differs, at least in degree, from the substitution of security parameters threat below, as Web Content Zones can control whether or not, for example, the browser executes unsigned active content.)

6.1.1.2 Delivery of Bad Accounting Information

In the case of CONTENT with value, CLIENTs may be inappropriately charged for viewing content that they did not successfully access. Conversely, some PUBLISHERs may reward CLIENTs for viewing certain CONTENT (e.g. programs that "pay" users to surf the Web). Should a CN fail to deliver appropriate accounting information, the CLIENT may not receive appropriate credit for viewing the required CONTENT.

Green, et. al.	Expires August 31, 2001	[Page 28]
Internet-Draft	CDI Architecture	March 2001

6.1.1.3 Delivery of Bad CONTENT

A CN that does not deliver the appropriate CONTENT may provide the user misleading information (either maliciously or inadvertently). This threat can be manifested as a failure of either the DISTRIBUTION SYSTEM (inappropriate content delivered to appropriate SURROGATES) or REQUEST-ROUTING SYSTEM (request routing to inappropriate SURROGATES, even though they may have appropriate CONTENT), or both. A REQUEST-ROUTING SYSTEM may also fail by forwarding the CLIENT request when no forwarding is appropriate, or by failing to forward the CLIENT request when forwarding is appropriate.

6.1.1.4 Denial of Service

A CN that does not forward the CLIENT appropriately may deny the CLIENT access to CONTENT.

6.1.1.5 Exposure of Private Information

CNs may inadvertently or maliciously expose private information (passwords, buying patterns, page views, credit card numbers) as it transits from SURROGATES to ORIGINS and/or PUBLISHERs.

6.1.1.6 Substitution of Security Parameters

If a SURROGATE does not duplicate completely the security facilities of the ORIGIN (e.g. encryption algorithms, key lengths, certificate authorities) CONTENT delivered through the SURROGATE may be less secure than the CLIENT expects.

6.1.1.7 Substitution of Security Policies

If a SURROGATE does not employ the same security policies and procedures as the ORIGIN, the CLIENT's private information may be treated with less care than the CLIENT expects. For example, the operator of a SURROGATE may not have as rigorous protection for the CLIENT's password as does the operator of the ORIGIN server. This threat may also manifest itself if the legal jurisdiction of the SURROGATE differs from that of the ORIGIN, should, for example, legal differences between the jurisdictions require or permit different treatment of the CLIENT's private information.

6.1.2 Threats to the PUBLISHER

6.1.2.1 Delivery of Bad Accounting Information

If a CN does not deliver accurate accounting information, the PUBLISHER may be unable to charge CLIENTs for accessing CONTENT or

Green, et. al.	Expires August 31, 2001	[Page 29]
Internet-Draft	CDI Architecture	March 2001

it may reward CLIENTs inappropriately. Inaccurate accounting information may also cause a PUBLISHER to pay for services (e.g. content distribution) that were not actually rendered.) Invalid accounting information may also effect PUBLISHERs indirectly by, for example, undercounting the number of site visitors (and, thus, reducing the PUBLISHER's advertising revenue).

6.1.2.2 Denial of Service

A CN that does not distribute CONTENT appropriately may deny CLIENTs access to CONTENT.

6.1.2.3 Substitution of Security Parameters

If a SURROGATE does not duplicate completely the security services of the ORIGIN (e.g. encryption algorithms, key lengths, certificate authorities, client authentication) CONTENT stored on the SURROGATE may be less secure than the PUBLISHER prefers.

6.1.2.4 Substitution of Security Policies

If a SURROGATE does not employ the same security policies and procedures as the ORIGIN, the CONTENT may be treated with less care than the PUBLISHER expects. This threat may also manifest itself if the legal jurisdiction of the SURROGATE differs from that of the ORIGIN, should, for example, legal differences between the jurisdictions require or permit different treatment of the CONTENT.

6.1.3 Threats to a CN

6.1.3.1 Bad Accounting Information

If a CN is unable to collect or receive accurate accounting information, it may be unable to collect compensation for its services from PUBLISHERs.

6.1.3.2 Denial of Service

Misuse of a CN may make that CN's facilities unavailable, or available only at reduced functionality, to legitimate customers or the CN provider itself. Denial of service attacks can be targeted at a CN's ACCOUNTING SYSTEM, DISTRIBUTION SYSTEM, or REQUEST-ROUTING SYSTEM.

Green, et. al. Expires August 31, 2001 [Page 30]

Internet-Draft CDI Architecture March 2001

6.1.3.3 Transitive Threats

To the extent that a CN acts as either a CLIENT or a PUBLISHER (such as, for example, in transitive implementations) such a CN may be exposed to any or all of the threats described above for both roles.

Green, et. al. Expires August 31, 2001 [Page 31]

Internet-Draft CDI Architecture March 2001

7. Acknowledgements

The authors would like to acknowledge the contributions and comments of Mark Day (Cisco), Fred Douglass (AT&T), Patrik Falstrom (Cisco), Don Gilletti (CacheFlow), Barron Housel (Cisco) John Martin (Network Appliance), Raj Nair (Cisco), Hilarie Orman (Novell), Doug Potter

Green, et. al.	Expires August 31, 2001	[Page 32]
Internet-Draft	CDI Architecture	March 2001

References

- [1] Hawkinson, J. and T. Bates, "Guidelines for creation, selection, and registration of an Autonomous System (AS)", BCP 6, March 1996,
<URL:<http://www.rfc-editor.org/rfc/bcp/bcp6.txt>>.
- [2] Postel, J., "Internet Protocol, DARPA Internet Program Protocol Specification", RFC 791, September 1981,
<URL:<http://www.rfc-editor.org/rfc/rfc791.txt>>.
- [3] Hares, S. and D. Katz, "Administrative Domains and Routing Domains A Model for Routing in the Internet", RFC 1136, December 1989,
<URL:<http://www.rfc-editor.org/rfc/rfc1136.txt>>.
- [4] Postel, J., "Domain Name Structure and Delegation", RFC 1591, March 1994,
<URL:<http://www.rfc-editor.org/rfc/rfc1591.txt>>.
- [5] Rekhter, Y. and T. Li, "A Border Gateway Protocol 4 (BGP-4)", RFC 1771, March 1995,
<URL:<http://www.rfc-editor.org/rfc/rfc1771.txt>>.
- [6] Carpenter, B., "Architecture Principles of the Internet", RFC 1958, June 1996,
<URL:<http://www.rfc-editor.org/rfc/rfc1958.txt>>.
- [7] Schulzrinne, H., Rao, A. and R. Lanphier, "Real Time Streaming Protocol", RFC 2326, April 1998,
<URL:<http://www.rfc-editor.org/rfc/rfc2326.txt>>.
- [8] Berners-Lee, T., Fielding, R. and L. Masinter, "Uniform Resource Identifiers (URI): Generic Syntax", RFC 2396, August 1998,
<URL:<http://www.rfc-editor.org/rfc/rfc2396.txt>>.
- [9] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P. and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2616, June 1999,
<URL:<http://www.rfc-editor.org/rfc/rfc2616.txt>>.
- [10] Carpenter, B., "Internet Transparency", RFC 2775, February 2000,
<URL:<http://www.rfc-editor.org/rfc/rfc2775.txt>>.
- [11] Cooper, I., Melve, I. and G. Tomlinson, "Internet Web

Green, et. al. Expires August 31, 2001 [Page 33]
Internet-Draft CDI Architecture March 2001

- [12] Volbrecht, J., Calhoun, P., Farrell, S., Gommans, L., Gross, G., de Bruijn, B., de Laat, C., Holdrege, M. and D. Spence, "AAA Authorization Framework", draft-ietf-aaa-authz-arch-00.txt (work in progress), October 1999,
<URL:http://www.ietf.org/internet-drafts/draft-ietf-aaa-authz-arch-00.txt>.
- [13] Day, M., Cain, B., Tomlinson, G. and P. Rzewski, "A Model for Content Internetworking", draft-day-cdn-model-05.txt (work in progress), March 2001,
<URL:http://www.ietf.org/internet-drafts/draft-day-cdn-model-05.txt>.
- [14] Day, M., Gilletti, D. and P. Rzewski, "Content Internetworking Scenarios", draft-day-cdn-scenarios-03.txt (work in progress), March 2001,
<URL:http://www.ietf.org/internet-drafts/draft-day-cdn-scenarios-03.txt>.
- [15] Gilletti, D., Nair, R., Scharber, J. and J. Guha, "Content Internetworking Authentication, Authorization, and Accounting Requirements", draft-gilletti-cdn-aaa-reqs-01.txt (work in progress), January 2001,
<URL:http://www.ietf.org/internet-drafts/draft-gilletti-cdn-aaa-reqs-01.txt>.
- [16] Barbir, A., Cain, B., Douglass, F., Green, M., Hofmann, M., Nair, R., Potter, D. and O. Spatscheck, "Known CDN Request-Routing Mechanisms", draft-cain-cdn-known-request-routing-01.txt (work in progress), February 2001.
- [17] Cain, B., Spatscheck, O., May, M. and A. Barbir, "Request-Routing Requirements for Content Internetworking", draft-ietf-cain-request-routing-req-01.txt (work in progress), March 2001.
- [18] Amini, L., Thomas, S. and O. Spatscheck, "Distribution Peering Requirements for Content Distribution Internetworking", draft-amini-cdi-distribution-reqs-00.txt (work in progress), February 2001.

Green, et. al. Expires August 31, 2001 [Page 34]
Internet-Draft CDI Architecture March 2001

Authors' Addresses

Mark Green
CacheFlow Inc.
650 Almanor Avenue
Sunnyvale, CA 94086
US

Phone: +1 408 543 0470
EMail: markg@cacheflow.com

Brad Cain
Cereva Networks

EMail: bcain@cereva.com

Gary Tomlinson
CacheFlow Inc.
12034 134th Ct. NE
Suite 201
Redmond, WA 98052
US

Phone: +1 425 820 3009
EMail: garyt@cacheflow.com

Stephen Thomas
TransNexus, Inc.
430 Tenth Street NW

Suite N204
Atlanta, GA 30318
US

Phone: +1 404 872 4887
EMail: stephen.thomas@transnexus.com

Phil Rzewski
Inktomi
4100 East Third Avenue
MS FC1-4
Foster City, CA 94404
US

Phone: +1 650 653 2487
EMail: philr@inktom.com

Green, et. al.	Expires August 31, 2001	[Page 35]
Internet-Draft	CDI Architecture	March 2001

Full Copyright Statement

Copyright (C) The Internet Society (2001). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC editor function is currently provided by the Internet Society.

Green, et. al.	Expires August 31, 2001	[Page 36]
----------------	-------------------------	-----------