Network Working Group
Internet-Draft
Expires: August 22, 2002

M. Day
Cisco
D. Gilletti
CacheFlow
P. Rzewski
Inktomi

February 22, 2002

Content Internetworking (CDI) Scenarios
draft-ietf-cdi-scenarios-00.txt

Status of this Memo

Copyright Notice

Abstract

In describing content internetworking as a technology targeted for use in the "real world", it's useful to provide examples of the possible sequence of events that may occur when two content networks decide to interconnect. The scenarios presented here seek to provide some concrete examples of what content internetworking is, and also to provide a basis for evaluating content internetworking proposals.

Table of Contents

1. Introduction

In [1], the concept of a "content network" is introduced and described. In addition to describing some general types of content networks, it also describes motivations for allowing content networks to interconnect (defined as ôcontent internetworkingö).

In describing content internetworking as a technology targeted for use in the "real world", it's useful to provide examples of the possible sequence of events that may occur when two content networks decide to interconnect. Naturally, different types of content networks may be created due to different business motivations, and so many combinations are likely.

This document first provides detailed examples of special cases of content networks that are specifically designed to participate in content internetworking (Section 2). We then discuss the steps that would be taken in order to "bring up" or "tear down" a content internetworking arrangement (Section 3). Next we provide some detailed examples of how content networks (such as those from Section 2) could interconnect (Section 4). Finally, we describe any security considerations that arise specifically from the examples presented here (Section 5).

The scenarios presented here answer two distinct needs:

1. To provide some concrete examples of what content internetworking is, and

2. To provide a basis for evaluating content internetworking proposals.

For details on the architectural framework used in the development of actual content internetworking protocols and interfaces, refer to [2]. For specific examples of systems where content internetworking has been implemented, refer to [5].

1.1 Terminology

Terms in ALL CAPS are defined in [1].

2. Special Cases of Content Networks

A CN is defined in [2] as having REQUEST-ROUTING, DISTRIBUTION, and ACCOUNTING interfaces. However, some participating networks may gravitate toward particular subsets of the CONTENT INTERNETWORKING interfaces. Others may be seen differently in terms of how they relate to their CLIENT bases. This section describes these refined cases of the general CN case so they may be available for easier

reference in the further development of CONTENT INTERNETWORKING scenarios. The special cases described are the Publishing Content

Network, the Brokering Content Network, and the Local Request-Routing Content Network.

2.1 Publishing Content Network

A Publishing Content Network (PCN), maintained by a PUBLISHER, contains an ORIGIN and has a NEGOTIATED RELATIONSHIP with two or more CNs. A PCN may contain SURROGATES for the benefit of serving some CONTENT REQUESTS locally, but does not intend to allow its SURROGATES to serve CONTENT on behalf of other PUBLISHERS.

Several implications follow from knowing that a particular CN is a PCN. First, the PCN contains the AUTHORITATIVE REQUEST-ROUTING SYSTEM for the PUBLISHER's CONTENT. This arrangement allows the PUBLISHER to determine the distribution of CONTENT REQUESTS among ENLISTED CNs. Second, it implies that the PCN need only participate in a subset of CONTENT INTERNETWORKING. For example, a PCN's DISTRIBUTION INTERNETWORKING SYSTEM need only be able to receive DISTRIBUTION ADVERTISEMENTS, it need not send them. Similarly, a PCN's REQUEST-ROUTING INTERNETWORKING SYSTEM has no reason to send AREA ADVERTISEMENTS. Finally, a PCN's ACCOUNTING INTERNETWORKING SYSTEM need only be able to receive ACCOUNTING data, it need not send it.

2.2 Brokering Content Network

A Brokering Content Network (BCN) is a network that does not operate its own SURROGATES. Instead, a BCN operates only CIGs as a service on behalf other CNs. A BCN may therefore be regarded as a "clearinghouse" for CONTENT INTERNETWORKING information.

For example, a BCN may choose to participate in DISTRIBUTION INTERNETWORKING and/or REQUEST-ROUTING INTERNETWORKING in order to aggregate ADVERTISEMENTS from one set of CNs into a single update stream for the benefit of other CNs. To name a single specific example, a BCN could aggregate CONTENT SIGNALS from CNs that represent PUBLISHERS into a single update stream for the benefit of CNs that contain SURROGATES. A BCN may also choose to participate in ACCOUNTING INTERNETWORKING in order to aggregate utilization data from several CNs into combined reports for CNs that represent PUBLISHERS.

This definition of a BCN implies that a BCN's CIGs would implement the sending and/or receiving of any combination of ADVERTISEMENTS and ACCOUNTING data as is necessary to provide desired services to other CONTENT NETWORKS. For example, a BCN only interested in aggregating ACCOUNTING data on behalf of other CNs would only need to have an ACCOUNTING INTERNETWORKING interface on its CIGs.

2.3 Local Request-Routing Content Network

Another type of CN is the Local Request-Routing CONTENT NETWORK (LCN). An LCN is defined as a type of network where CLIENTS' CONTENT REQUESTS are always handled by some local SERVER (such as a caching proxy [1]). In this context, "local" is taken to mean that both the CLIENT and SERVER are within the same administrative domain, and there is an administrative motivation for forcing the local mapping. This type of arrangement is common in enterprises where all CONTENT REQUESTS must be directed through a local SERVER for access control purposes.

As implied by the name, the LCN creates an exception to the rule that there is a single AUTHORITATIVE REQUEST-ROUTING SYSTEM for a particular item of CONTENT. By directing CONTENT REQUESTS through the local SERVER, CONTENT RESPONSES may be given to CLIENTS without first referring to the AUTHORITATIVE REQUEST-ROUTING SYSTEM. Knowing this to be true, other CNs may seek a NEGOTIATED RELATIONSHIP with an LCN in order to perform DISTRIBUTION into the LCN and receive ACCOUNTING data from it. Note that once it's participating in DISTRIBUTION INTERNETWORKING and ACCOUNTING INTERNETWORKING, the SERVERS within the LCN effectively take on the role of SURROGATES. However, an LCN would not intend to allow its SURROGATES to be accessed by non-local CLIENTS.

This set of assumptions implies multiple things about the LCN's CONTENT INTERNETWORKING relationships. First, it is implied that the LCN's DISTRIBUTION INTERNETWORKING SYSTEM need only be able to send DISTRIBUTION ADVERTISEMENTS, it need not receive them. Second, it is implied that an LCN's ACCOUNTING INTERNETWORKING SYSTEM need only be able to send ACCOUNTING data, it need not receive it. Finally, due to the locally defined REQUEST-ROUTING, the LCN would not participate in REQUEST-ROUTING INTERNETWORKING.

3. Content Internetworking Arrangements

When the controlling interests of two CNs decide to interconnect their respective networks (such as for business reasons), it is expected that multiple steps would need to occur.

The first step would be the creation of a NEGOTIATED RELATIONSHIP. This relationship would most likely take the form of a legal document that describes the services to be provided, cost of services, SLAs, and other stipulations. For example, if an ORIGINATING CN wished to leverage another CN's reach into a particular country, this would be laid out in the NEGOTIATED RELATIONSHIP.

The next step would be to configure CONTENT INTERNETWORKING protocols on the CIGs of the respective CNs in order to technically support the terms of the NEGOTIATED RELATIONSHIP. To follow our

previous example, this could include the configuration of the ENLISTED CN's CIGs in a particular country to send DISTRIBUTION ADVERTISEMENTS to the CIGs of the ORIGINATING CN. In order to

configure these protocols, technical details (such as CIG addresses/hostnames and authentication information) would be exchanged by administrators of the respective CNs.

In the event that the controlling interests of two CNs no longer wish to have their networks interconnected, it is expected that these tasks would be undone in reverse order. That is, first the protocol configurations would be changed to cease the movement of ADVERTISEMENTS and/or ACCOUNTING data between the networks. After this, the NEGOTIATED RELATIONSHIP would be legally terminated.

4. Content Internetworking Scenarios

This section provides several scenarios that may arise in CONTENT INTERNETWORKING implementations.

Note that we obviously cannot examine every single permutation. Specifically, it should be noted that:

- Any one of the interconnected CNs may have other CONTENT INTERNETWORKING arrangements that may or may not be transitive to the relationships being described in the diagram.

- The graphical figures do not illustrate the CONTENT REQUEST paths. It is assumed that the direction of CONTENT REQUESTS follow the methodology given in [2] and that the end result is that a REQUEST-ROUTING SYSTEM eventually returns to the CLIENT the IP address of the SURROGATE deemed appropriate to honor the CLIENT's CONTENT REQUEST.

The scenarios described include a general case, two cases in which BCNs provide limited interfaces, a case in which a PCN enlists the services of multiple CNs, and a case in which multiple CNs enlist the services of an LCN.

4.1 General Content Internetworking

This scenario considers the general case where two or more existing CNs wish to establish a CONTENT INTERNETWORKING relationship in order to provide increased scale and reach for their existing customers. It assumes that all of these CNs already provide REQUEST-ROUTING, DISTRIBUTION, and ACCOUNTING services and that they will continue to provide these services to existing customers as well as offering them to other CNs.

In this scenario, these CIs would interconnect with others via a CIG which provides a REQUEST-ROUTING INTERNETWORKING SYSTEM, a DISTRIBUTION INTERNETWORKING SYSTEM, and an ACCOUNTING INTERNETWORKING SYSTEM. The net result of this interconnection would be that a larger set of SURROGATES will now be available to the

CLIENTS.

FIGURE 1 shows three CNs which have interconnected to provide greater scale and reach to their existing customers. They are all participating in DISTRIBUTION INTERNETWORKING, REQUEST-ROUTING INTERNETWORKING, and ACCOUNTING INTERNETWORKING.

As a result of the NEGOTIATED RELATIONSHIPS it is assumed that:

1. CONTENT that has been INJECTED into any one of these ORIGINATING CNs may be distributed into any other ENLISTED CN.

2. Commands affecting the DISTRIBUTION of CONTENT may be issued within the ORIGINATING CN, or may also be issued within the ENLISTED CN.

3. ACCOUNTING information regarding CLIENT access and/or DISTRIBUTION actions will be made available to the ORIGINATING CN by the ENLISTED CN.

4. The ORIGINATING CN would provide this ACCOUNTING information to the PUBLISHER based on existing Service Level Agreements (SLAs).

5. CONTENT REQUESTS by CLIENTS may be directed to SURROGATES within any of the ENLISTED CNs.

The decision of where to direct an individual CONTENT REQUEST may be dependent upon the DISTRIBUTION and REQUEST-ROUTING policies associated with the CONTENT being requested as well as the specific algorithms and methods used for directing these requests. For example, a REQUEST-ROUTING policy for a piece of CONTENT may indicate multiple versions exist based on the spoken language of a CLIENT. Therefore, the REQUEST-ROUTING SYSTEM of an ENLISTED CN would likely direct a CONTENT REQUEST to a SURROGATE known to be holding a version of CONTENT of a language that matches that of a CLIENT.

FIGURE 1 - General CONTENT INTERNETWORKING

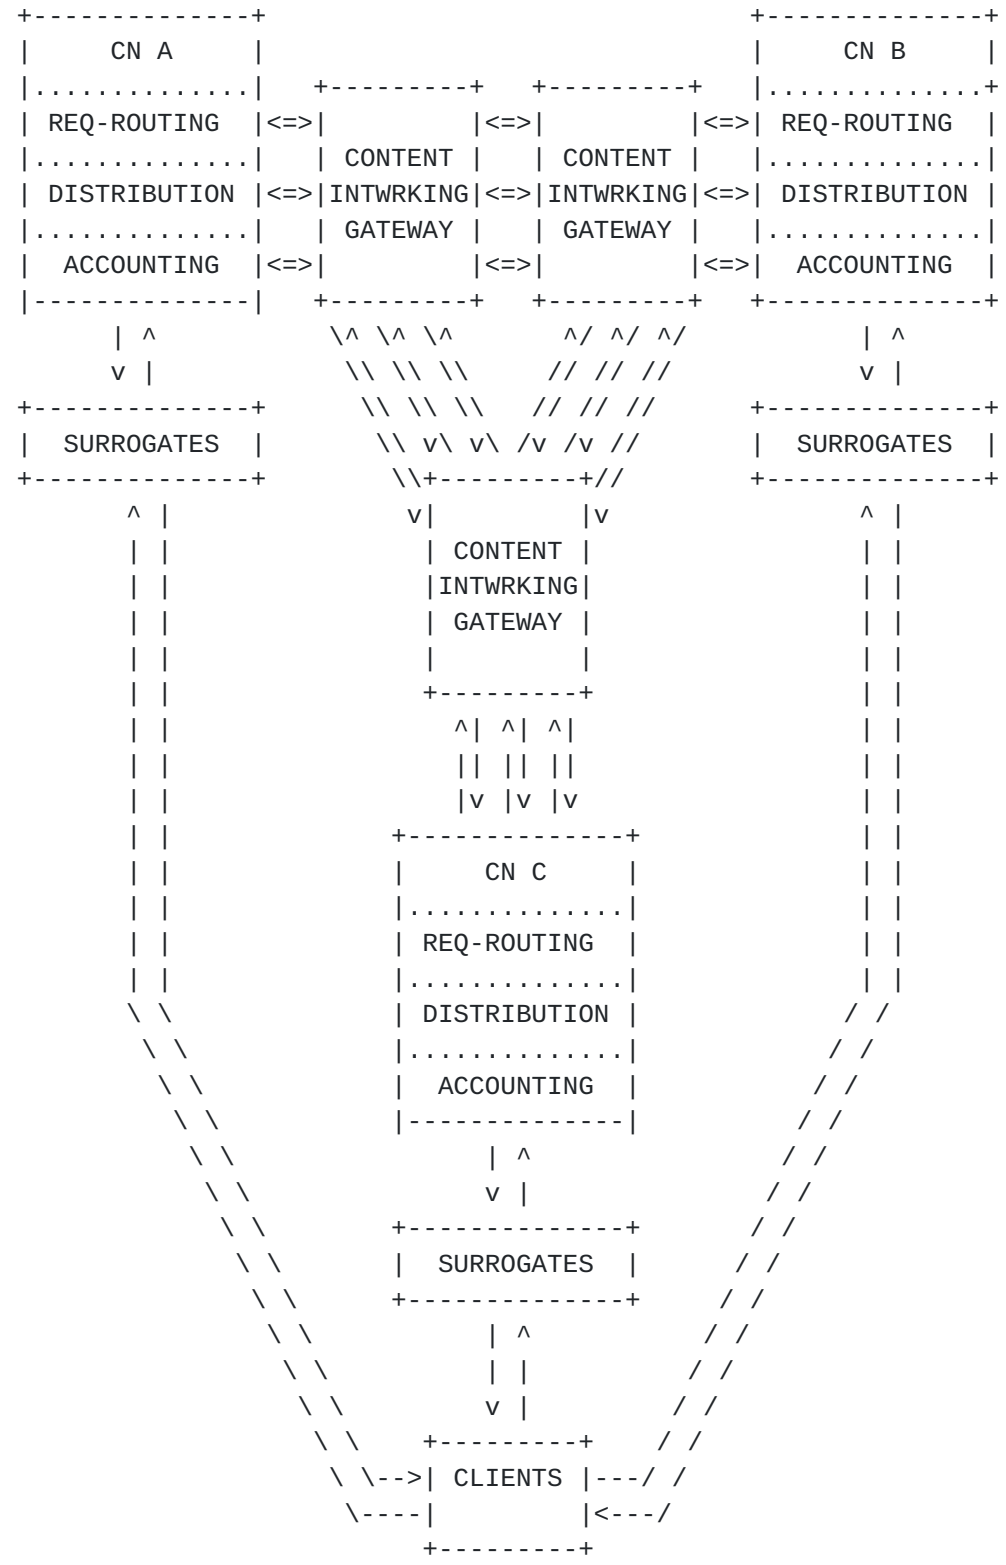

4.2 BCN providing ACCOUNTING INTERNETWORKING and REQUEST-ROUTING INTERNETWORKING

This scenario describes the case where a single entity (BCN A) performs ACCOUNTING INTERNETWORKING and REQUEST-ROUTING INTERNETWORKING functions, but has no inherent DISTRIBUTION or DELIVERY capabilities. A potential configuration which illustrates this concept is given in FIGURE 2.

In the scenario shown in FIGURE 2, BCN A is responsible for collecting ACCOUNTING information from multiple CONTENT NETWORKS (CN A and CN B) to provide a clearinghouse/settlement function, as well as providing a REQUEST-ROUTING service for CN A and CN B.

In this scenario, CONTENT is injected into either CN A or CN B and its DISTRIBUTION between these CNs is controlled via the DISTRIBUTION INTERNETWORKING SYSTEMS within the CIGs. The REQUEST-ROUTING SYSTEM provided by BCN A is informed of the ability to serve a piece of CONTENT from a particular CONTENT NETWORK by the REQUEST-ROUTING SYSTEMS within the interconnected CIGs.

BCN A collects statistics and usage information via the ACCOUNTING INTERNETWORKING SYSTEM and disseminates that information to CN A and CN B as appropriate.

As illustrated in FIGURE 2, there are separate REQUEST-ROUTING SYSTEMS employed within CN A and CN B. If the REQUEST-ROUTING SYSTEM provided by BCN A is the AUTHORITATIVE REQUEST-ROUTING SYSTEM for a given piece of CONTENT this is not a problem. However, each individual CN may also provide the AUTHORITATIVE REQUEST-ROUTING SYSTEM for some portion of its PUBLISHER customers. In this case care must be taken to ensure that the there is one and only one AUTHORITATIVE REQUEST-ROUTING SYSTEM identified for each given CONTENT object.

FIGURE 2 - BCN providing ACCOUNTING INTERNETWORKING and REQUEST-ROUTING INTERNETWORKING

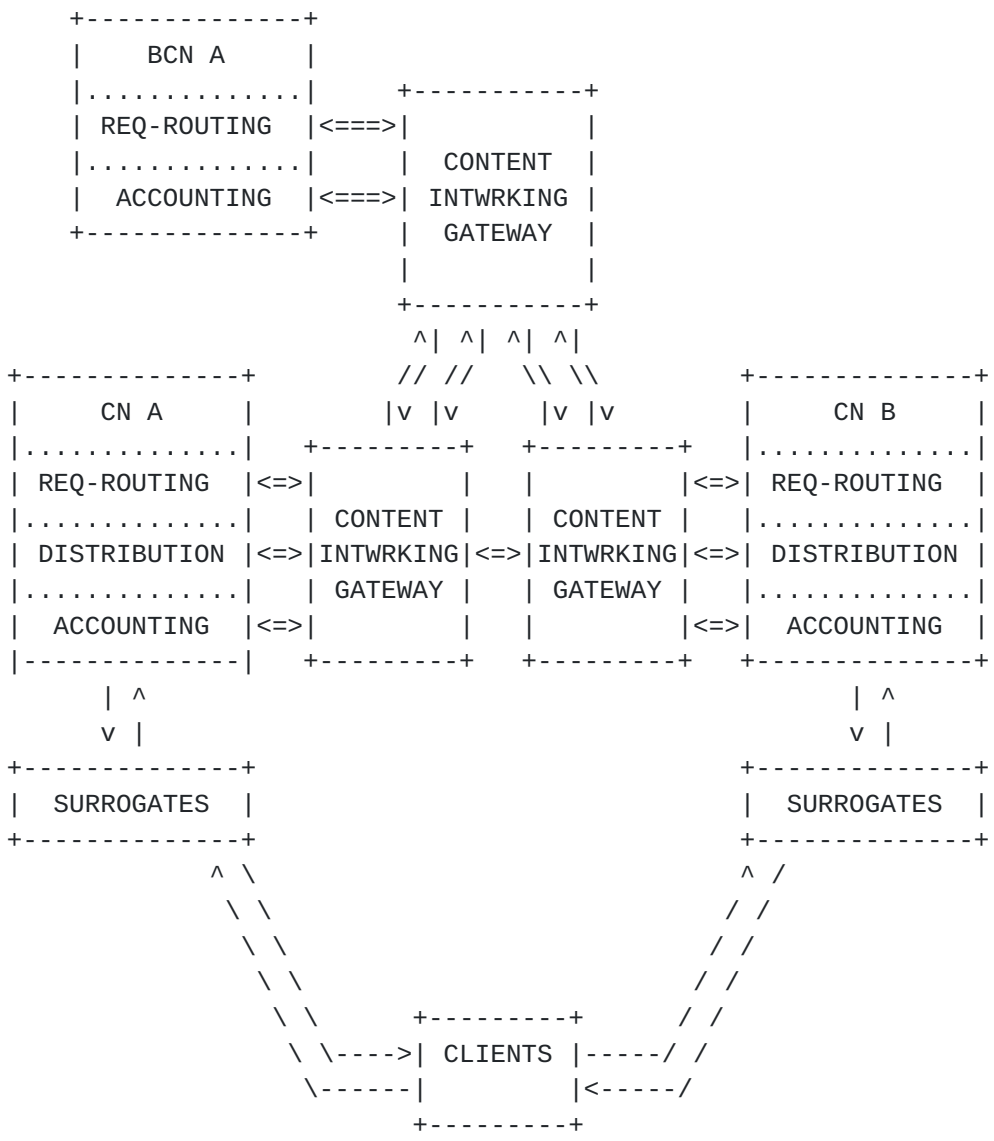

4.3 BCN providing ACCOUNTING INTERNETWORKING

This scenario describes the case where a single entity (BCN A) performs ACCOUNTING INTERNETWORKING to provide a clearinghouse/ settlement function only. In this scenario, BCN A would enter into NEGOTIATED RELATIONSHIPS with multiple CNs that each perform their own DISTRIBUTION INTERNETOWRKING and REQUEST-ROUTING INTERNETWORKING as shown in FIGURE 3.

FIGURE 3 - BCN providing ACCOUNTING INTERNETWORKING

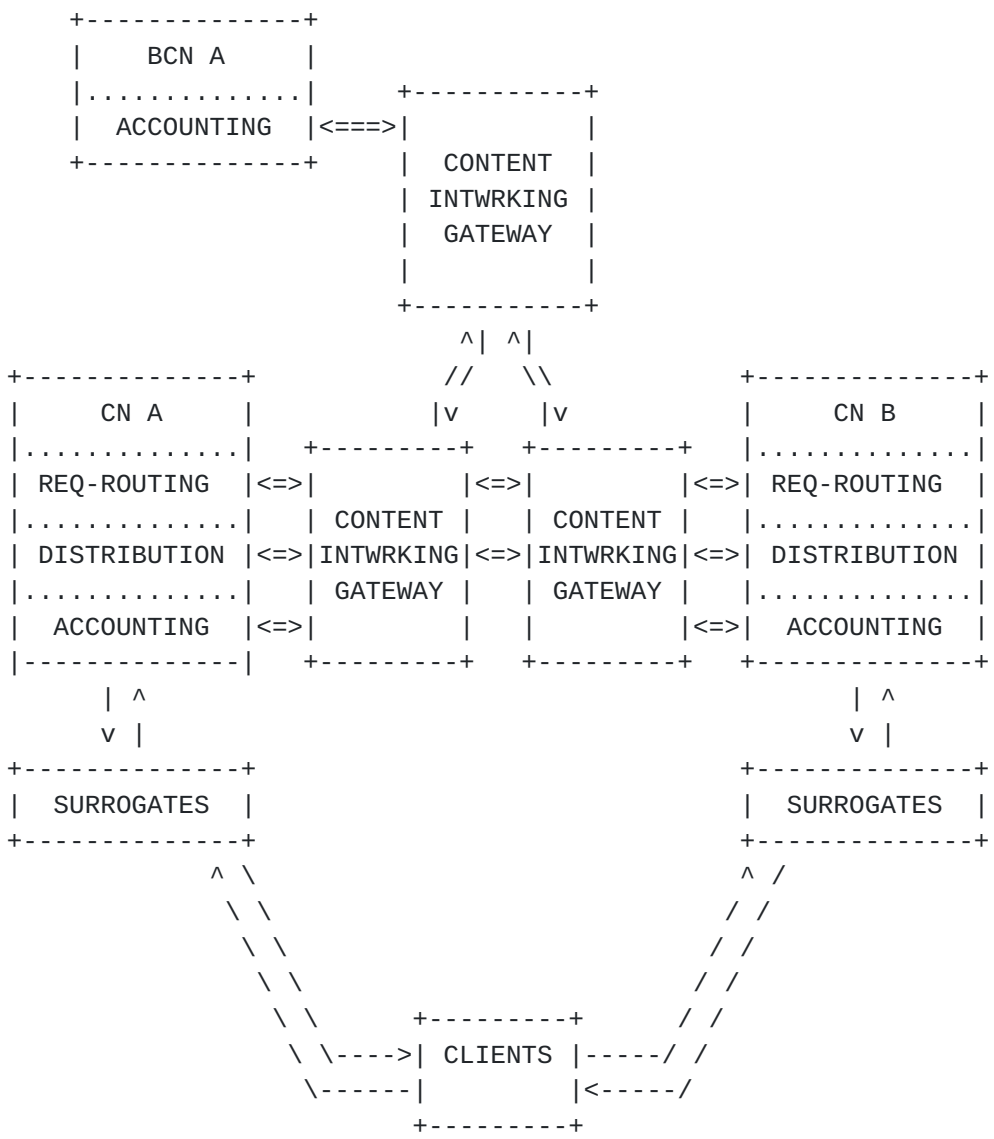

4.4 PCN ENLISTS multiple CNs

In the previously enumerated scenarios, PUBLISHERS have not been discussed. Much of the time, it is assumed that the PUBLISHERS will allow CNs to act on their behalf. For example, a PUBLISHER may designate a particular CN to be the AUTHORITATIVE REQUEST-ROUTING SYSTEM for its CONTENT. Similarly, a PUBLISHER may rely on a particular CN to aggregate all its ACCOUNTING data, even though that data may originate at SURROGATES in multiple distant CNs. Finally, a PUBLISHER may INJECT content only into a single CN and rely on that CN to ENLIST other CNs to obtain scale and reach.

However, a PUBLISHER may wish to maintain more control and take on the task of ENLISTING CNs itself, therefore acting as a PCN (Section 2.1). This scenario, shown in FIGURE 4, describes the case where a PCN wishes to directly enter into NEGOTIATED RELATIONSHIPS with multiple CNs. In this scenario, the PCN would operate its own CIG and enter into DISTRIBUTION INTERNETWORKING, ACCOUNTING INTERNETWORKING, and REQUEST-ROUTING INTERNETWORKING relationships with two or more CNs.

FIGURE 4 - PCN ENLISTS multiple CNs

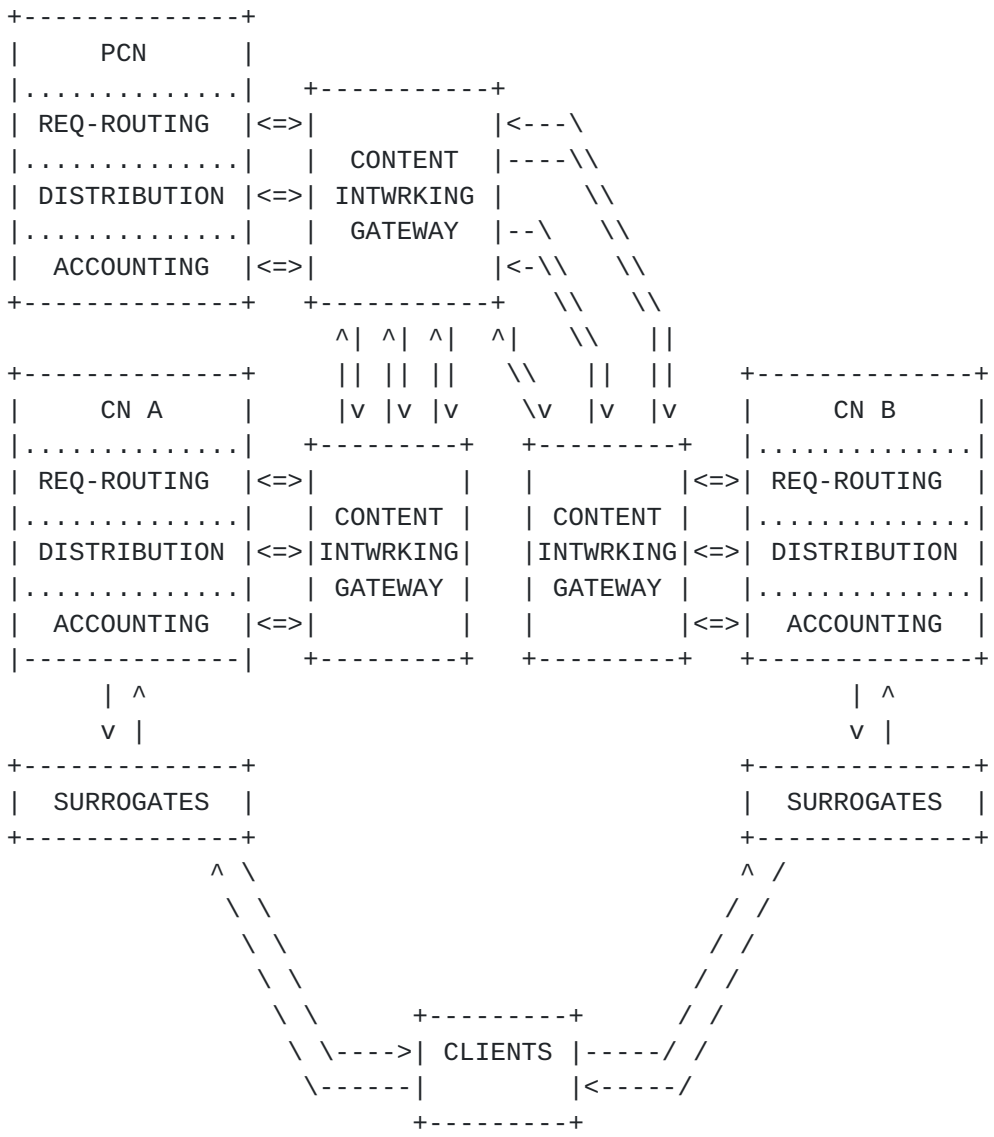

4.5 Multiple CNs ENLIST LCN

A type of CN described in Section 2.3 is the LCN. In this scenario, we imagine a tightly administered CN (such as within an enterprise) has determined that all CONTENT REQUESTS from CLIENTS must be serviced locally. Likely due to a large CLIENT base in the LCN, multiple CNs determine they would like to engage in DISTRIBUTION INTERNETWORKING with the LCN in order to extend control over CONTENT objects held in the LCN's SURROGATES. Similarly, the CNs would like to engage in ACCOUNTING INTERNETWORKING with the LCN in order to receive ACCOUTING data regarding the usage of the content in the local SURROGATES. This scenario is shown in FIGURE 5.

FIGURE 5 - Multiple CNs ENLIST LCN

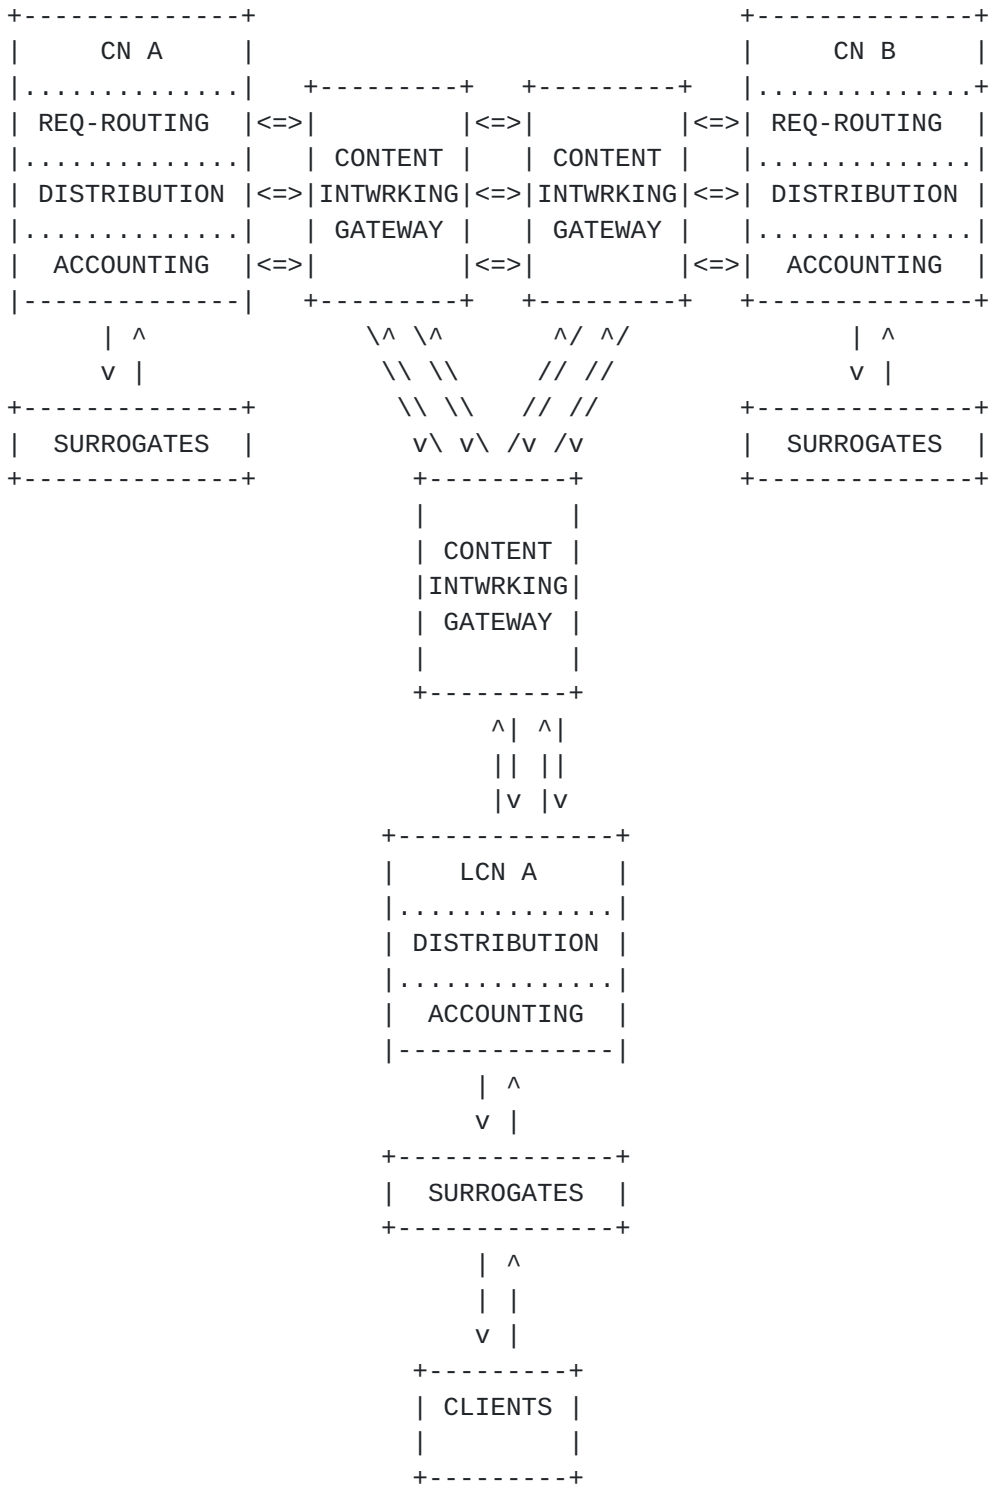

5. Security Considerations

This section contains security considerations that arise specifically from the examples presented here. For a more general discussion of security in the CDI protocols, see [2].

Due to the likely reliance on ACCOUNTING data as the basis of payment for services, the likelihood of fraud may be a concern of parties that participate in CONTENT INTERNETWORKING. Indeed, it's easy to imagine fabricating log entries or increasing throughput numbers to increase revenue. While this is a difficult problem to solve, there are some approaches to be explored. A useful tool would be a "fraud detection" analysis tool that is capable of modeling human usage patterns and detecting anomalies. It may be logical for such a tool to be run by a BCN that is acting as an "impartial third party", ENLISTED only to ensure fairness among participants. Additionally, a BCN may be ENLISTED to perform random audits of ACCOUNTING data.

6. Acknowledgements

The authors acknowledge the contributions and comments of Fred Douglis (AT&T), Raj Nair (Cisco), Gary Tomlinson (CacheFlow), John Scharber (CacheFlow), Nalin Mistry (Nortel), Steve Rudkin (BT), Christian Hoertnagl (IBM), Christian Langkamp (Oxford University), and Don Estberg (Activate).

References

[1] Day, M., Cain, B., Tomlinson, G., and P. Rzewski, "A Model for Content Internetworking (CDI)", draft-ietf-cdi-model-00.txt (work in progress), February 2002, <URL:http://www.ietf.org/internet-drafts/draft-ietf-cdi-model-00.txt>.

[2] Green, M., Cain, B., Tomlinson, G., Thomas, S., and P. Rzewski, "Content Internetworking Architectural Overview", draft-ietf-cdi-architecture-00.txt (work in progress), February 2002, <URL:http://www.ietf.org/internet-drafts/draft-ietf-cdi-architecture-00.txt>.

[3] Gilletti, D., Nair, R., Scharber, J., and J. Guha, "CDN-I Internetworking Authentication, Authorization, and Accounting Requirements", draft-ietf-cdi-aaa-reqs-00.txt (work in progress), February 2002, <URL:http://www.ietf.org/internet-drafts/draft-ietf-cdi-aaa-reqs-00.txt>.

[4] Aboba, B., Arkko, J. and D. Harrington, "Introduction to

Accounting Management", RFC 2975, October 2000,

<URL:ftp://ftp.isi.edu/in-notes/rfc2975.txt>.

[5] Douglis, F., Chaudhri, I. and P. Rzewski, "Known Mechanisms for Content Internetworking", draft-douglis-cdi-known-mech-00.txt, November 2001, <URL:http://www.ietf.org/internet-drafts/draft-douglis-cdi-known-mech-00.txt>.

Authors' Addresses

Mark S. Day
Cisco Systems
135 Beaver Street
Waltham, MA 02452
US

Phone: +1 781 663 8310
EMail: markday@cisco.com

Don Gilletti
CacheFlow, Inc.
441 Moffett Park Drive
Sunnyvale, CA 94089 USA
US

Phone: +1 408 543 0437
EMail: don@cacheflow.com

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

Phone +1 650 653 2487
Email: philr@inktomi.com

Full Copyright Statement

document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other

Acknowledgement

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