Network Working Group
Internet-Draft
Expires: December 26, 2001

D. Gilletti
CacheFlow
R. Nair
Cisco
J. Scharber
CacheFlow
J. Guha
Apogee
June 2001

Content Internetworking (CDI) Authentication, Authorization, and Accounting Requirements

draft-gilletti-cdnp-aaa-reqs-01.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 December 26, 2001.

Copyright Notice

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

Abstract

In developing a solution for CDN Internetworking it is necessary to define and accommodate requirements for Authentication, Authorization, and Accounting. Since the Authentication, Authorization, and Accounting (AAA) working group is already focused on defining these requirements, this document attempts to leverage that work. It contains the requirements that would have to be supported by a AAA service to formulate a solution for CDN peering.

Table of Contents

1. Introduction

The initial model for the World Wide Web (WWW) was based on clients interacting with origin servers to request and receive content or services. As the Web increased in scale this model proved unwieldy for several reasons and resulted in current industry efforts to build and operate Content Distribution Networks or CDNs. The overall purpose of these CDNs is to create a scalable service that can meet aggregate client demand while improving the performance and quality of delivery. With increased demand for CDN services a need has been generated for a mechanism for interconnecting or peering these systems.

A typical CDN has relationships with publishers and provides them with accounting and access related information. This information is typically provided in the form of aggregate or detailed log files.

In addition, these CDNs typically collect accounting information to aid in operation, billing and SLA verification. Since all accounting data is collected within the CDN's administrative domain there is no requirement for generalized systems or protocols.

Peering or interconnecting these CDNs introduces the need to obtain similar accounting data from a foreign domain. This requirement means that customers of a peered CDN service (publishers, clients, and CDNs) must now have a generalized or standard means of obtaining accounting information to support current as well as planned business models. For example, the desire to implement business models such as "Pay Per View" may require that there exist a mechanism for authenticating and authorizing clients at a delivery point that lies in a foreign domain.

This document along with [4],[5], and [6] outline requirements to be satisfied in order to develop a mechanism for interconnecting or peering CDNs. The intent of this set of documents is to provide structure and guidelines for the evaluation of proposed solutions.

This document is focused on describing requirements for the Accounting Peering System as described in [6]

This document frames the requirements for the Accounting Peering System against the ongoing work of the AAA working group. This was done because the authors realized that considerable effort has already been expended in identifying inter-domain trust models and accounting requirements within that working group. Therefore, a conscious decision has been made to leverage that existing body of work before making additional proposals. As such, this document relies heavily on RFC 2904[1], RFC 2975[2], and RFC 2977[3].

Since the concentration of this effort is to determine the requirements for CDN internetworking / peering, the accounting requirements within an individual CDN are largely ignored within this document.

The core actions and activities within the CDN-I domain is essentially enumerated as Content Injection, Content Distribution, Content Request, Content/Service Delivery, and Content Retrieval. These are the primary activities that need to be tracked and accounted. Please refer architectural diagrams in [4],[5], and [6].

This document focuses and details the requirements on the digital representation of the above activities, along with the means and mechanisms to exchange these representations between peering parties which have participated in the activity, and / or any other component in a secure and guaranteed manner.

Requirements for the remaining CDI architectural elements, the Request Routing System, which is responsible for directing user agents to the distributed content, and the Distribution Peering Requirements for CDI, which is responsible for distributing content between a CDN and the elements that it, are detailed in [9], [10].

1.1 Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC-2119 [8].

2. Terminology

This section introduces new terminology not already defined in RFC 2975[2], RFC 2977[3], or [4].

CDN Service:
An action that is directly or indirectly related to the act of moving content from publisher to consumer.

Customer:
A billable entity (typically a publisher, client or peered CDN) that agrees to exchange compensation for a CDN Service.

Entitlement:
A right to access a given CDN Service or content object typically provided to a customer by a provider.

Flat Rate:
Indicates there is no limit on the amount of CDN Service that a customer can consume during a Period.

Percentile:
Indicates that a CDN Service will be billed at a rate that is based on a multiplier (Usage Rate) times the Usage during the Period.

Period:
The duration for which the Usage counter or Entitlement is active.

Provider:
An entity that offers a CDN Service in exchange for Compensation.

Unit Of Measure (UOM):
Indicates how Usage should be tracked (i.e. minutes, seconds, bytes, etc).

Usage:
A counter that measures the access or use of a CDN Service by the Customer.

Usage Rate:
A per-unit cost associated with the Usage of a CDN Service.

Pricing / Rating Tiers:
Indicates the existence of a schedule against which Usage of a given CDN Service is tracked and billed.

Work :
This is the definition of an activity in which a CDN partakes by providing a specific function, role or service. These activities are restricted to only those which involve the participation of peering entities outside the CDNs administrative domain.

Tier :
This is the conceptual enclosure of a specific function (such as accounting, settlement, rating, billing, invoicing, payment, provisioning etc) in the context of a layered / phased multi-function environment.

CDR (Content Detail Record):
This is the digital entity which capture the 'what', 'who', 'when', 'where', and 'how' of the work done.

3. Overview of Accounting Peering System

The Accounting Peering system is responsible for the definition, generation and exchange of usage / consumption data entities (CDRs) which depict the activities and works performed, requested, and completed between peering CDNs, and any component internetworking with a CDN. These CDRs typically will contain 'what work was done', 'who did the work', 'when was the work done', 'how was the work done', 'Resources Used' etc.

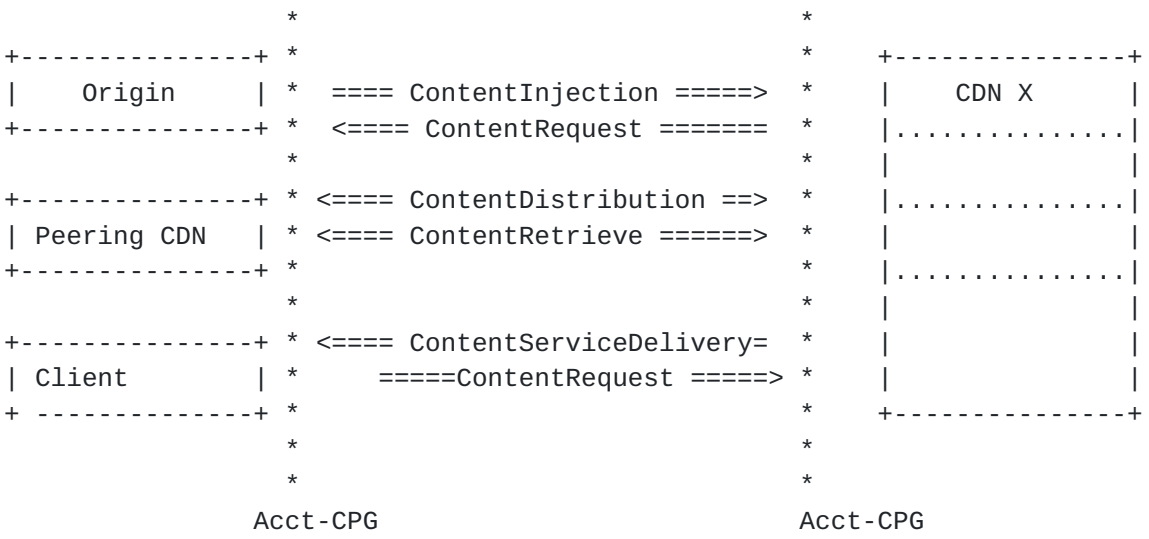

Figure 1 : Accounting Peering System Components and Activities

Fig 1 illustrates the architectural entities which will internetwork with a CDN as well as the activities that transpire between any two peering entities . Each activity is to be considered as discrete accountable events in their own right. The arrows indicate the parties and roles involved in each activity as each activity has a source and destination role in the communications exchange (note : the arrows do not indicate who is generating / receiving the CDR)

The model above allows for complex chaining and sequencing of activities which express the relationship of interoperations that a typical CDN will encounter.

The core activites (as described above) that need to be accounted are ContentInjection, ContentDistribution, ContentRequest, ContentRetrieval, and ContentServiceDelivery. The activities above cannot span multiple internetworking hops.

Each independent piece of activity or work performed by / to a CDN, can be transmitted and collected by any entity or instrumentation, which implements the Accounting-Peering System Interface(Acct-CPG).

It is envisioned that each CDN would have an instrumentation platform which would detect or be informed of the activities described above. This platform or detection mechanism is not in-scope of Acct-CPG.

A transport protocol would be used to exchange CDRs between an entity which generates a CDR, and an entity which receives a CDR. Presumably, the entity which is capable of detecting the activity will be the entity that generates the CDR to an entity interested in receiving the CDR. The Acct-CPG interface will not specify which entity can, must, or should implement the Acct-CPG interface.

Accounting systems in general will have to support the real-time occurrence of the above mentioned activities. That is when the activities do take place, the activity must be recorded and not lost. Reliability of recording accountable events / activities is required or otherwise the accounting infrastructure will be compromised. However the Acct-CPG will not specify the mechanism to record the activities that have occurred, but will solely impress upon the necessity of reliability and integrity in the process of activity detection, and recording.

The nature of certain activities may also span a segment of time from activity initiation to completion. Acct-CPG systems shall be able to generate interim, and composite CDRs of each discrete activity which depict the passage of time and states of an activity. Thus, the Acct-CPG interface shall be charged with the responsibility of support for interim and composite CDRs, and support for real time and/or offline/batch exchange of accountable works / activities.

The Acccounting data entities (CDRs) may be used by other downstream functional tiers such as Rating & Billing, Capacity Planning, Performance & Monitoring Analysis, Payment systems, Account Management, Settlement Systems etc. These downstream tiers are considered out-of-scope.

4. Assumptions

Certain assumptions / expectation of the operating environment are detailed below. Should any of the assumptions change, there may be material implications on the requirements of the CDN-I Accounting Peering (Acct-CPG) systems.

4.1 Firewalls

There are no firewalls between the path of the Accounting Peering Systems. Peering CDN-I Acccounting systems can establish a communication channel between themselves provided they have the appropriate and valid trust credentials.

4.2 CDR Generation & Reception

Any entity (network element or service element) can be a CDR Generator and/or a CDR Receiver as long as the entity is 'CDN-I Acct-CPG' enabled.

4.3 Storage Requirements

The CDN-I Accounting Peering requirements shall not place any constraints or restrictions on the storage of CDRs. CDR Storage is essentially out of scope. However, to provide failover and recovery in the data exchange protocol, there most likely will be storage implications for an entity to be considered 'CDN-I Accounting Peering' enabled.

4.4 Authentication & Authorization

The CDN-I Accounting Peering requirements shall only exchange Authentication & Authorization Messages to enable the exchange of CDRs. The policies, and mechanisms which influence these Authentication and Authorizations messages are to be provided and maintained by an external functional component or process, and are to be considered out-of-scope, except to the extent that it influences the requirements of accounting data exchange.

4.5 Measurement / Metering Systems

Instrumentation (hard and/or soft) exists on the CDN-I network which can detect, recognize, and inform an Accounting-Peering System when a ContentInjection, ContentDistribution, ContentRequest, ContentRetrieval, and ContentServiceDelivery act / event has taken place.

4.6 Inter-Domain Trust

All systems are conformant with the AAA Authentication and Authorization Framework for Inter-Domain trust. Refer [1]

4.7 Proxy CDN-I Account-Peering systems

If Accounting-Peering (Acct-CPG) system communicate through intermediate Acct-CPG systems, it is necessary that the CDR payload is secure between the originator Acct-CPG and target Acct-CPG server. The security

requirement may be end-to-end security, CDR payload integrity, confidentiality, replay protection, and non-repudiation. Refer [1].

5. Accounting 'Works'

The objective of this section is to define the requirement of a scaleable framework for depicting all the various acts / services / 'works' performed by any entity which internetworks with a CDN, and which needs to be accounted. It is envisioned that a finite (core) set of works will need to be defined so as to achieve initial adoption and traction. Thereafter the framework must accommodate expansion of acts / services / work in the CDN Internetworking domain.

5.1 Introduction

An expression of 'work' performed consists of a collection/sequence of attributes which consist of identifiers, measures, and counters. The attributes serve to answer the who, what, where, and how of these elements of usage or acts of consumption (CDRs).

For the CDN peering/interconnection scenario, it is important to construct the expressions/acts of consumption such that there is no dispute and ambiguity in meaning between parties producing and receiving these expressions. In each act of consumption, there is a minimum set of attributes in which an act/expression of consumption/usage is considered "complete and undisputable" and therefore able to be used by any downstream tiers (ex : rating and billing).

5.2 General Requirements

5.2.1 Framework

A framework which can accommodate the scaleable definition of CDRs is required. This framework shall be able to introduce new CDRs and their associated schemas at a later timeframe.

The framework must ensure that 'Context' of the CDR payload is available to any party such that domains / scope of CDR and attribute applicability / validity is defined. 'Context' will mean roles of participating entities, domain, and scope. There shall be no overloading of attribute meanings. Every measure and counter must have have units, minima and maxima, and incrementals defined. Every identifier must be persistent within a defined domain and timeframe. Decisions on in-band or out-of-band context embedding will be needed and shall be defined in the framework. The framework shall support self-description of the usage attributes.

This framework shall NOT define any actions, rules, constraints and context with regards to the processing of CDRs.

XML technology should be considered as a vehicle to achieve the above framework. Other strategies can also be considered if the end-value

is achievable.

5.2.2 Form-Factor of CDRs

The representation and form-factor of the CDRs must also be addressed. Binary based and character based form factors need to be supported.

5.2.3 Timing Requirement of CDR exchange

Certain CDRTypes will have delivery requirements which meet timing constraints. For each independent operating scenario and 'work' (CDR) type, requirements on CDR transport exchange and timings need to be assessed. It will be required to specify timing constraints for certain CDR Types, and will be used by the data exchange protocol during the channel setup phase.

5.2.4 Identifiers

An identifier is an attribute which associates a name convention to an entity. Examples of identifies are OrganizationName, EndUserName, Name of Movie, ContentID etc These identifiers can and do have properties and characteristics such as persistence, scope & domain, time-to-live etc. In the accounting context, these identifiers must be resolvable, unambiguous, recoverable and unique in the applicable domain.

5.2.5 Measures

Measures are attributes that help to describe 'how' a certain piece of work was performed, delivered or consumed. Typical examples are QoS, JitterDelay, etc. Measures need to be defined completely and without ambiguity by defining known minima, maxima, unit of increments etc. The definition of measures must remain persistent and consistent within its scope and domain.

5.2.6 Counters

Counters are attributes that help to describe the quantity of resources consumed by a singular accountable act / work. Typically Counters must be defined by its units, minima and maxima and the mathematical operations that are permissible on a counter.

5.2.7 Name Space Convention (Distributed,Domained,Global)

In a distributed environment, a domain consistent way of naming entities such as a customer, ContentID, ContentType etc is required, so that any member participating in the CDN Internetworking activities can be consistently identified. The scope / domain of applicability of every identifier must be referenceable by any party participating in the work represented by the specific CDR, and / or by any party involved in the exchange of the specific CDR.

Likewise, the framework must also provision a Name Space mechanism for

naming a specific content within the federation. For example, a specific movie or song might need to be named in the same way in all of the different points of the federation, independently of the domain that is delivering the specific content. This SHOULD follow a standard that can be interpreted by every member of the federation.

5.2.8 Security Requirements

This document assumes that the solutions suggested within this document will be compliant with the trust model given in RFC 2904[1].

5.3 Core Works Requirements

This section defines those 'work' items that form the initial core set of 'works' that must be supported by any 'CDN-I Acccounting' enabled entity. The objective here is to achieve consensus around a set of accountable works which are deemed sufficiently important such that its definition and support is warranted.

5.3.1 Introduction

For each of the accountable works defined below in the subsections, the following must be defined :

5.3.1.1 List of Attributes

Each attribute must define its name, its meaning, whether it is a required / optional / conditional attribute, its data type (int, string, complex etc)

5.3.1.2 Encoding and Representation

The encoding and representation format of each attribute inside the CDR must be defined.

5.3.1.3 Use Case Model

Generally describes the involved entities in the creation of the CDR. The 'Context' of the CDR will most likely be defined here as well.

5.3.1.4 Work Flow Model

Details the sequencing of events of the involved entities which generates the CDR and where the CDR has to be sent.

5.3.1.5 Work State Model

Details the state transition diagram of the 'Work' by defining the triggers and the state of work from inception to completion. The 'State' of Work may or may not be distributed across one or more elements in the federation.

5.3.2 ContentInjection Work

This 'work' is the event(CDR) that represents the act of a Host Origin Server which successfully 'injects' a piece of 'content' into a CDN for further distribution. An example scenario is where a Content Provider or a Publisher wishes to distribute content. The Publisher typically transfer the relevant content to a CDN Service Provider. This transaction is referred to as Injection.

A logical representation of this CDR is as below :

| CDRType | TimeStamp | ContentID | ContentType | OriginID | CDN_ID | ContentByteSize | State | ErrorCode |

Attribute Name	Type	Description
CDRType	String	Type of CDR (Value = ContentInjection)
TimeStamp	Long	Time of Content Injection transaction (Start/End)
ContentURI	String	Unique identifier for the injected content
ContentType	String	Type of Content (WebPage, Movie, Song etc)
Origin ID	String	Unique Identifier for Content Source
CDN_ID	String	Unique Identifier for CDN ServiceProvider
ContentByteSize	Long	Size of Content Injected in Bytes
State	String	State of Injection (start,complete,error)
ErrorCode	String	UnAuthorized,UnAuthenticated,TimeOut,etc

5.3.3 ContentDistribution Work

This 'work' is the event (CDR) that represents the act where a CDN 'distributes' the 'content' to another peering CDN.

A logical representation of this CDR is as below :

| CDRType | TimeStamp | Content_ID | ContentType | CDNSrcID | CDNDestID| | ContentByteSize | State | ErrorCode |

Attribute Name	Type	Description
CDRType	String	Type of CDR (Value = ContentDistribution)
TimeStamp	Long	Time of Content Distribution transaction (start/end)
ContentURI	String	Unique identifier of content to be distributed
ContentType	String	Type of Content (WebPage, Movie, Song etc)
CDNSrcID	String	Unique Identifier of src distributing content
CDNDestID	String	Unique Identifier of dest receiving content
ContentByteSize	Long	Size of Content distributed in Bytes
State	String	Distribution State (start,complete,error)
ErrorCode	String	UnAuthorized,UnAuthenticated,TimeOut,etc

5.3.4 ContentRequest Work

This transaction is the event (CDR) that represents the act where an end-user (EU) request access to a specific Content.

A logical representation of this CDR is as below :

| CDRType | TimeStamp | ContentURI | ContentType | ContentByteSize |
| RequesterID | ReceiverID | State | ErrorCode |

Attribute Name	Type	Description
CDRType	String	Type of CDR (Value = ContentRequest)
TimeStamp	Long	Time of content request transaction (start/end)
ContentURI	String	Unique identifier for the content to be distributed
ContentType	String	Type of Content (WebPage, Movie, Song etc)
RequesterID	String	Unique Identifier for entity requesting the content
ReceiverID	String	Unique Identifier for entity receiving request
ContentByteSize	Long	Size of Content distributed in Bytes
State	String	State of Request (start,complete,error)
ErrorCode	String	UnAuthorized,UnAuthenticated,TimeOut,etc

5.3.5 ContentRetrieval Work

This transaction is the event (CDR) that represents the act where when a Content Request 'Miss' occurs, the 'content' is retrieved from the origin server and delivered to the element (CDN / cache) where the miss occurred.

A logical representation of this CDR is as below :

| CDRType | TimeStamp | ContentURI | ContentType | ContentByteSize |
| Origin_ID | Receiver_ID | State | ErrorCode |

Attribute Name	Type	Description
CDRType	String	Type of CDR (Value = ContentRetrieval)
TimeStamp	Long	Time of Content Retrieval transaction (start/end)
ContentURI	String	Unique identifier for the retrieved content
ContentType	String	Type of Content (WebPage, Movie, Song etc)
Origin ID	String	Unique Identifier for Content Source
Receiver_ID	String	Unique Identifier for receiver of content retrieved
ContentByteSize	Long	Size of Content Injected in Bytes
State	String	State of Injection (start,complete,error)
ErrorCode	String	UnAuthorized,UnAuthenticated,TimeOut,etc

5.3.6 ContentServiceDelivery Work

This transaction is the event (CDR) that represents the act where a CDN delivers the Content requested to the entity which requested the content.

| CDRType | TimeStamp | ContentServiceURI | ContentServiceType |
| ContentByteSize | DeliverерID | ReceiverID | State | ErrorCode |

Attribute Name	Type	Description
CDRType	String	Type of CDR (Value = Content/Service Delivery)
TimeStamp	Long	Time of Content Delivery transaction (start/end)
Content/ServiceURI	String	Unique identifier the delivered content/service
Content/ServiceType	String	Type of Content/Service (WebPage, Movie, Song etc)
DeliveryMode	String	mode of delivery (stream, filetransfer, http, etc)
DelivererID	String	Unique Identifier of entity delivering the content
ReceiverID	String	Unique Identifier for entity receiving the content
ContentByteSize	Long	Size of Content Injected in Bytes
State	String	State of Delivery (start,complete,error)
ErrorCode	String	UnAuthorized,UnAuthenticated,TimeOut,etc

[editors note : Measures and counters which express how the content / service was delivered is conditionally dependant on the content/service type and delivery mode (ex : videostream (frames-per-sec, MPEG-level), Songs (AudioQoS),]

6. Data Exchange Mechanism / Protocol

6.1 Introduction

This objective of this section is to develop the requirements of a transport mechanism which shall be responsible for the transfer / exchange of a 'Work/Activity' (CDR) from one entity to another entity.

6.2 General Requirements

This section details some of the general requirements of a transport protocol.

6.2.1 Separation of Exchange Protocol from CDR payload

The transfer protocol must be cleanly decoupled from the CDR payloads that it will transfer.

6.2.2 Transfer Capability Negotiation

Support for push, pull and poll transfers modes need to be supported.

6.2.3 Singular, Batched, Flow, & RealTime Modes of Data Exchange

The transport protocol shall support batch, flow, and realtime modes of exchange of CDR payloads. Support for multiple channels of transport must exist to accommodate multiple varying throughput rate requirements, and/or multiple exchange modes between the accounting peering parties which occur at the same time. A specific transport channel shall be able to exchange multiple CDRs of a singular CDRType. That is, mixed CDRTypes within a singular channel is not supported.

6.2.4 Efficient Encoding

6.2.5 Transfer Flow & Rate Control

The transfer protocol shall support flow control mechanisms to achieve sustainable delivery throughput between the two data exchange peers.

6.2.6 Guaranteed Delivery

It is to be assumed that all works that are to be accounted for MUST never be lost. Therefore all transfer modes must achieve reliable and guaranteed delivery of CDR payloads. Unless there is a compelling case for an unguaranteed delivery requirement, this assumption and requirement shall stand.

6.2.7 CDR Relationship Identifiers

CDR EventIdentifiers (unique) may be required if relationships exist across a set of CDRs. Situations where interim CDRs are generated, it is necessary to track the sequencing of a related set to ensure completeness, detect errors, and the retransmission of CDRs. If relationships of a set of CDR spans a distributed domain, then a distributed numbering strategy must exist.

6.2.8 Protocol State Machines

Clear visualization of transport protocol state machines in the sender and receiver must be developed.

6.2.9 Encryption

It is recommended that encryption works from other IETF standards be leveraged to ensure data / CDR security.

6.3 Authorization and Authentication

Authorization and authentication mechanisms must exist in the data exchange protocol to enable the initiation of data exchange. This mechanism may be influenced by external (out-of-scope) policy and control mechanisms / processes which precede the transfer / exchange of CDRs.

6.4 CDR Receivers

Any entity that is 'CDN-I Accounting' enabled is eligible to receive a CDR. To enable the delivery of CDRs, the CDR Receiver must contact a CDR Generator and establish a channel. A channel must only be able to transfer CDRs of a singular CDRType.

6.5 CDR Generators :

Only 1 CDR must be generated for a singular incidence of 'work'. Any entity which is 'CDN-I Accounting' enabled, is eligible to generate the CDR. The CDR Generator entity must be able to support the delivery of a CDR stream to multiple recipients if a single channel has been created between each recipient and the CDR Generator.

6.6 Fault-Tolerance

Fault Tolerance, failover, and recovery mechanism mechanisms are required to insure against network failure, accounting peering component failure, packet loss, and/or device reboots.

7. Further Issues

8. Recommendations

[Editor's Note: This section is here only to record some ideas that need to be discussed while this specification is being progressed.]

One means of accommodating these types of services is to build off of the ongoing work of the IETF AAA working group [2]. At present this work is centered on the DIAMETER framework and protocol suite for both provisioning and accounting. Early observations indicate that DIAMETER has several characteristics that are desirable for consideration in fulfilling these accounting requirements. The high point characteristics are that it:

Has a model that supports either direct aggregation to home provider 3rd party brokering.

Has well developed security and trust relationships.

Supports standardized, extensible accounting record format.

Is generally extensible via object oriented techniques.

The general model of extending DIAMETER is to define required extensions to the protocol much like one would do to an abstract base class in C++ via base class and subclassing.

Although its a bit premature to fully assess the suitability of DIAMETER to meet these requirements, early observations indicate that it sets forth a reasonable framework from which to develop a base model for this effort.

Early observations have also identified the following issues with the model that will likely create a need for the following extensions to the base framework:

1. DIAMETER works on a request-by-request basis like pay-per-view. While this model is okay for some applications, it will have to be extended to support cases where a CDN pays at a larger granularity (e.g., by a million content hits) and then resells to its users or another CDN. This would apply to cases where a CDN subscribes to a peered billing organization or pays for distribution in a peering CDN. Existing DIAMETER mechanisms could be used for pay-per-view content inside a CDN but may need a higher level protocol across CDNs for aggregate content programming. This protocol SHOULD co-exist with DIAMETER message proxying. It can borrow message routing models from DIAMETER (e.g. realm-based routing).

2. DIAMETER uses end-to-end security. This may not work well across

CDN boundaries. As previously discussed, it may be necessary to be flexible about the definition of the "end" to be the CDN boundary. This will be consistent with the need for CDNs to serve as a content provisioning entity and makes it possible to aggregate request traffic.

3. DIAMETER needs to be extended with AVPs specific to web-based billable events.

More detailed analysis needs to be undertaken before these conclusions can be validated.

9. Conclusion

There is a considerable amount of work remaining in defining the accounting requirements and relationships. As such, the authors welcome additional input from interested parties.

10. Acknowledgements

The authors acknowledge the contributions and comments of Brad Cain (Mirror Image), Mark Day (Cisco), Fred Douglis (AT&T), John Martin (Network Appliance), Doug Potter (Cisco), Oliver Spatscheck (AT&T), Gary Tomlinson (Entera), Lisa Amini (IBM) and Abhi Deskmukh (Apogee).

11 References

[1] Vollbrecht, J., Calhoun, P., Farrell, S., Gommans, L., Gross, G., de Bruijn, B., de Laat, C., Holdrege, M. and D. Spence, "AAA Authorization Framework", RFC 2904, August 2000, <ftp://ftp.isi.edu/in-notes/rfc2904.txt>.

[2] Aboba, B., Arkko, J. and D. Harrington, "Introduction to Accounting Management", RFC 2975, October 2000, <ftp://ftp.isi.edu/in-notes/rfc2975.txt>.

[3] Glass, S., Hiller, T., Jacobs, S. and C. Perkins, "Mobile IP Authentication, Authorization, and Accounting Requirements", RFC 2977, October 2000, <ftp://ftp.isi.edu/in-notes/rfc2977.txt>.

[4] Day, M., Cain, B. and G. Tomlinson, "A Model for CDN Peering", draft-day-cdnp-model-03.txt, (work in progress), November 2000, <http://www.ietf.org/internet-drafts/draft-day-cdnp-model-03.txt> .

[5] Day, M. and D. Gilletti, "CDN Peering Scenarios", draft-day-cdnp-scenarios-02.txt, (work in progress), November 2000, <http://www.ietf.org/internet-drafts/draft-day-cdnp-scenarios-02.txt>

[6] Green, M., Cain, B. and G. Tomlinson, "CDN Peering Architectural Overview", draft-green-cdnp-framework-00.txt, (work in progress), September 2000, <http://www.ietf.org/internet-drafts/draft-green-cdnp-gen-arch-02.txt>

[7] IPDR NDM 2.0 'Network Data Management - Usage for IP Services' <http://www.ipdr.org>

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

[9] Cain, B., Spatscheck, O., May, M. and A. Barbir, "Request Routing Requirements for Content Internetworking", January 2001.

[10] Amini, Thomas S., Spatscheck, O., "Distribution Peering Requirements for Content Distribution Internetworking", January 2001.

Authors' Addresses

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

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

Raj Nair
Cisco Systems
50 Nagog Park
Acton, MA 01720
US

Phone: +1 978 206 3029
EMail: rnair@cisco.com

John Scharber
CacheFlow, Inc.
441 Moffett Park Drive
Sunnyvale, CA 94089
US

Phone: ???
EMail: john.scharber@cacheflow.com

Jay Guha
Apogee Networks
Park 80 West, Plaza II,
Saddle Brook, NJ
Tel: +1 201 368 8800
Email: jayg@apogeenet.com

Full Copyright Statement

Copyright (C) The Internet Society (2000). 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.