



**B R O A D B A N D  
C O N T E N T  
D E L I V E R Y F O R U M**

## Requirements for End-To-End Delivery of Broadband Content

Oct. 22, 2001

Version 1.0

THIS WORK PRODUCT IS RELEASED “AS IS” AND WITHOUT ANY WARRANTY OF ANY KIND, INCLUDING, WITHOUT LIMITATION, ANY EXPRESS OR IMPLIED WARRANTY OF TITLE, NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. IN NO EVENT SHALL BROADBAND CONTENT DELIVERY FORUM, INC. OR ANY OF ITS MEMBERS BE LIABLE FOR ANY INDIRECT, SPECIAL, EXEMPLARY, INCIDENTAL OR CONSEQUENTIAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

COPYRIGHT © 2001 by Broadband Content Delivery Forum, Inc. All rights reserved.  
2570 W. El Camino Real, Ste. 304, Mountain View, CA 94040, USA, 650/949-6730.

## **1 Abstract**

The mission of the Broadband Content Delivery Forum is to bring together the leading players in the services/applications and infrastructure arena to develop and foster the conditions and standards necessary for industry participants to create, promote and deliver a plethora of network-delivered, affordable, value-added services to residential and business users. Two working groups have been established within the forum, one focusing on technical issues and one focused on business issues. One goal of the BCD Forum is to understand the infrastructure requirements for supporting the end-to-end setup and delivery of broadband content from the viewpoint of content providers, content network service providers, network providers, and content users. In addition, by understanding the needed requirements, the BCD Forum will promote current standards and influence future standards that support these requirements. This document was created by the combined efforts of both technologists and content providers and enumerates a set of requirements that need to be satisfied to support the broadband “experience.” Although the basic framework in the document addresses the requirements for many different vertical applications, the initial main focus is on personalized, client/server applications for real-time delivery of audiovisual content. Specific requirements for other types of applications, such as peer-to-peer applications will be addressed in later versions of this or other documents. This “living” document will be maintained and updated on a continual basis by the BCD Forum and on behalf of the content and service provider communities.

<b>1</b>	<b>ABSTRACT</b>	<b>2</b>
<b>2</b>	<b>INTRODUCTION</b>	<b>5</b>
2.1	Content Value Chain	6
2.2	Network Supported Applications	7
2.3	Document Focus	10
2.4	Requirements Categories	10
2.5	Vocabulary	11
<b>3</b>	<b>FOUNDATION REQUIREMENTS</b>	<b>12</b>
<b>4</b>	<b>OFFLINE DISTRIBUTION OF CONTENT AND DESCRIPTIONS</b>	<b>14</b>
4.1	Content Provider	14
4.2	Content Network	15
4.3	Content User	15
<b>5</b>	<b>LOCATING CONTENT</b>	<b>15</b>
5.1	Content Provider/Content Network	15
5.2	Network Provider	15
5.3	Content User	15
<b>6</b>	<b>NEGOTIATING THE DELIVERY OF CONTENT</b>	<b>16</b>
6.1	Content Provider/Content Network	17
6.2	Content Network/Network Provider	18
6.3	Network Provider	19
6.4	Content User	19

<b>7</b>	<b>DELIVERY OF CONTENT FROM SOURCE TO TARGET</b>	<b>20</b>
7.1	Content Provider	20
7.2	Content Network	20
7.3	Network Provider	21
7.4	Content User	22
<b>8</b>	<b>EXECUTION OF POST-DELIVERY MANAGEMENT</b>	<b>22</b>
8.1	Content Provider	23
8.2	Network Provider	23
8.3	Content Network	24
8.4	Content User	24
<b>9</b>	<b>AUTHORS ADDRESSES</b>	<b>25</b>
<b>10</b>	<b>ACKNOWLEDGEMENTS</b>	<b>25</b>
<b>11</b>	<b>REFERENCES</b>	<b>25</b>

## 2 Introduction

The BCD Forum has observed that there is a set of requirements that need to be fulfilled by a network infrastructure to support the end-to-end setup and delivery of broadband content, particularly audiovisual content. The goal of this document is to enumerate these requirements from the standpoint of four major groups who participate in this broadband “experience”.

The four groups with requirements are as follows:

- Content Provider – the person or organization that controls the content and receives accounting information concerning the content. This could be the original content creator (e.g. Disney, Time-Warner,) an agent or aggregator with permission to provide the content, or a user capable of creating broadband content. This has also been called the broadband content-owner. Note that we distinguish between the content creator and the content provider (although they could be the same organization or person). There may be business requirements between the creator and content providers and information may flow between the entities. The technologies for delivering the content between them, however, may be the same as characterized as between the content provider and content user. For example, just as a content provider delivers data to many content users, a content creator could distribute data to multiple content providers. The main requirement for content creators is that they receive the appropriate accounting and distribution information based on service level agreements that they may have in effect.
- Content Network - an organization that provides a network facility for the delivery, peering, security, and accounting of content. Content Networks include network infrastructure that exists in the OSI layers 4/7 and deals with the routing and forwarding of requests and responses for content. The units of transported data in Content Networks, such as web pages, movies, or songs, can be very large and may span hundreds or millions of packets. Content networks can also provide mediation for games and peer-to-peer delivery of content. Typically a content network consists of a request-routing system, surrogates (a content server other than the origin server) for replicating content on behalf of content providers, and an accounting system. (Definition from IETF Content Peering group)<sup>1</sup>
- Network Provider – an organization that provides a network facility for the delivery and accounting of data delivery. This lower layer (OSI layers 1-3) network infrastructures revolves around the routing, forwarding, and switching of frames and packets, There are two main types of network providers: *Core* providers, such as Qwest, transfer data within the core of the Internet; and *access* providers such as Verizon, provide access (DSL, Cable, Dial-up) to Content Users at the edge of the network. During the delivery of information between content provider and Content User, there may be one or two *access* providers (for any given Content User there will probably only be one *access* provider) and one or more *core* providers.

- Content User – the individual and his/her associated consumption devices wishing to consume the content (e.g. view it on a PC, watch it on TV, listen to it on a phone, etc.). Note that the requesting and consumption devices/agents need not be the same. For example, a WAP-enabled device with user profile and billing information could be used to order a movie to be delivered to a set-top box.

Note that this classification is categorized by function. Single organizations, however, may provide multiple functions. For example, an organization like Verizon could be both a content network (owns its own servers that store and deliver content) and a network provider (owns the access networks to end users).

## **2.1 Content Value Chain**

Content, created by the original content provider, can flow to other groups in the value chain and ultimately to the Content User. The possible flow of content is outlined in Table 1. In all scenarios, content flows from source to target via one or more network providers (these network providers may have different business infrastructures which need to be accommodated). For example, in the Syndication model, Disney may create original content but gives another organization, aka Intertainer, the rights to control and distribute the content. Intertainer may use a content network, aka Akamai, to push the content closer to Content Users, and so on. Each of these models is not mandatory, and the document does not exclude the simple case of a single content provider and Content User. The most complex model, however, must be considered when deriving the end-to-end requirements.

Function	Source	Target
Syndication	CP (Disney)	CP (Intertainer)
Replication	CP (Intertainer)	CN (Akamai)
Content Network Internetworking	CN (Akamai)	CN (Digital Island)
Client Server	CP (Yahoo)	CU(Fred)
Edge Service	CN (Akamai)	CU(Sue)
Peer to Peer	CU (Fred)	CU(Sue)

Table 1. Example scenarios for content transfer  
CP: Content Provider  
CN: Content Network  
CU: Content User

Note that we do not address the set of possible contractual relations (e.g. service level agreements, distribution rights, etc.) that may exist between the groups. These relations may, however, imply the flow of distribution rights, auditing information, etc. and may have explicit requirements.

## 2.2 Network Supported Applications

There are many different vertical applications that can be realized by a broadband network. These include, video-on-demand, music downloads, distributed games, etc. Understanding the relationship between applications and the network is crucial in determining what network equipment to build, what capabilities are needed and when they are needed. In addition, understanding an application’s behavioral attributes is important when setting up the delivery of content for a particular end-to-end experience. For example, knowing that an application needs high-bandwidth delivery and low delay jitter may impact the decision on which network can be used to transport the content. Several useful paradigms have been developed<sup>2,3</sup> that broadly categorize applications by the characteristics of the network services required for successful delivery of the application. Those described by Siegal<sup>2</sup> in “*Network Supported Applications*,” are:

- **Hosted Applications** - These are applications where content users initiate requests for content and content providers or content networks respond by pushing content to the user. This is the basis of most web-based or http applications. Hosted applications can involve streaming or reliable file transfers. Video-on-demand and digital music downloads are example applications in this space.
- **Broadcast Applications** – Broadcast Applications involve the simultaneous pushing of content to multiple recipients. There are two main distribution solutions for enabling broadcast applications in a network. In the network IP-multicast model, all routers and switches between source and destination must support and be enabled for IP-multicast. In the application multicast model, an overlay network is created using point-to-point connections to build a distribution tree. The near-video-on-demand model, where content providers periodically broadcast content at pre-defined times is an example application.
- **Peer-to-peer Applications** - Peer-to-Peer applications are those applications where the content is not under the direct control of the content provider or content network. Instead the Network Provider provides connectivity between content users. These applications uniquely allow people to communicate with people. Some current popular applications in this category include AOL Instant Messenger, Microsoft NetMeeting and Napster.

Table 2 provides an example list of possible applications in the three different categories.

Market	Hosted	Broadcast	Peer-to-Peer
Entertainment	Video-on-demand	Pay-per-view	Post-production collaboration
Medicine	Recorded Medical Procedures	Live Medical Procedures	Medical collaboration

Education	On-line courses	Course webcasts	Virtual classroom
Publishing	News/sports libraries	News/sports channels	Discussion chat rooms

Table 2. Example applications

In addition, each of these applications could be defined at varying bandwidths, with varying amounts of security, with different billing options, etc. Thus, a set of attributes for the applications has been identified and are as follows:

1. **Quality of Content and Speed of Delivery** - This attribute refers to the quality level of the content needed for this application. Quality levels are defined by comparison to different, more common media (e.g. “near VHS quality,” “CD quality”). Delivery speed required may be real-time ( i.e., the time it takes to experience time dependent media such as music or video ) or may be faster or slower.
2. **Real-Time** - A real-time media event is one that is being experienced as it is being streamed over the network. It may be live or recorded. If it is real-time, the total end-to-end delay that is acceptable is important. This would be a combination of the network delay and jitter buffer. The amount of packet loss that is acceptable is also important. Frequently, real-time applications do not allow for retransmission of lost, late or garbled packets.
3. **Reliability / Availability** - This measures the application’s availability and includes both network and application issues. Applications involving the transmission of live events would put the most stress on availability because other applications could be tried again at a later time.
4. **Sources / Receivers** - This refers to whether the application is point-to-point, point-to-multipoint, or multipoint-to-multipoint. Each of the target applications needed to support one or more of these distribution methods. Webcasting is point-to-multipoint and may have large numbers of receivers. eCinema would have fewer receivers, but larger bandwidth. Post-Production would need point-to-point and video collaboration would need multipoint-to-multipoint.
5. **Knowledge of Receivers** - This refers to whether point-to-multipoint distribution would be to open groups or closed groups. Open groups allow anyone to join a webcast, while closed groups make a webcast available only to a pre-defined group.
6. **Interactivity** - Interactivity can be with the application or with other users. Interactivity with users demands that the network comply with issues of human perception and expectations. Interactions with multimedia applications may be limited to choosing what is played and when it starts or it may involve greater interaction including VCR controls or even gaming.
7. **Security** - Security refers to those aspects that protect the content, the infrastructure, and information about content usage from attack and from

unauthorized access. Zones of Data Center security based on firewalls combined with encryption, secure sockets and VPNs are what is needed here. Issues of trust, denial of service and subversion should also be considered. For example, New Media frequently involves strong security issues.

8. **Digital Rights Management** - Digital Rights Management (DRM) is different from security in that it limits how an authorized receiver of the data may use it. This may include limits on storing, copying playing, etc. It may also include time dependent limits. Watermarking technology could be used to keep track of particular copies. Eventually, there may be a need for the network to do watermarking to indicate the path through the network any digital content file took. At first, however, DRM will be a client server issue rather than a network issue.
9. **Storage Capability** - This refers to how much media content must be available to the end-user. Storage is important both because some of the files needed for these applications are extremely large (measured in Terabytes) and because some applications demand large libraries of different files.
10. **Standards Compliance and Client Interworking** - Some of the standards compliance issues at the network level include TCP/IP, UDP, VPNs, IP Multicast and RTP/RTSP/RTCP. Media encoding standardized formats include the MPEG family. Popular, yet proprietary streaming formats include Real Networks formats and Microsoft Windows Media formats
11. **Application Capabilities / Features / Components** - This would include billing capabilities, menu generation, spot insertion, indexing of stored content, etc. For example, digital music downloads would need both menu generation and a way to bill by title.
12. **Application Management** - This includes any application specific provisioning or data collection. Closed group multicasting needs ways for people to sign up for services. This could be web-based, agent – assisted, or some combination of these. Bulk provisioning may be necessary as groups are defined and added to a closed group. Application management systems need to be fully integrated with other aspects of managing the networks, server environments and content in these systems.
13. **Symmetry** – Does the application require symmetric capabilities from the network? For symmetric applications, the upstream and downstream bandwidths to/from a content user would be the same. Thus, this attribute applies to applications where information and/or content flows both to and from a content user.

Based on the various types of applications and attributes, Table 3 presents several example scenarios. Due to space constraints, only a subset of the attributes is shown for each.

Application	Type	Speed of delivery	Real-Time	Sources Receivers	Interactivity	Symmetry
-------------	------	-------------------	-----------	-------------------	---------------	----------

Video Conferencing	Peer-to-peer	64-384 Kbps	Yes	Point-to-point and multipoint-to-multipoint	High	Required
Video-on-demand	Hosted	56Kbps-19Mbps	Yes	Point-to-point	Low (FF/Rew, etc.)	Not required
Pay-per-view	Broadcast	56Kbps-19Mbps	Yes	Point-to-multipoint	None	Not required
E-Learning Webcast	Broadcast	56Kbps-1.5Mbps	Yes	Point-to-multipoint	Low	Not required
Digital music downloads	Hosted	Best effort	No	Point-to-point	None	Not required
File Sharing (aka Naptster)	Peer-to-peer	Best effort	No	Point-to-point	None	Not required

Table 3. Examples of applications and their attributes

### 2.3 Document Focus

This document focuses on hosted applications for real-time delivery of audiovisual content. The basic framework in the document, however, addresses broadcast and peer-to-peer applications, albeit at a high-level. In addition, many of the requirements outlined are equally pertinent in the three application types described in the previous section. For example the, Single-Source Multicast initiative has a lot in common with the hosted applications model.

Details about broadcast and peer-to-peer applications have not been explored and need further development. For example, specific methods for multicast or peer-to-peer communication have not been enumerated. Likewise, symmetrical communication applications represent an important class of broadband services with unique requirements, also well worth further treatment.

Lastly, although this document focuses on “broadband content delivery”, many of the requirements are needed in narrowband environments as well.

### 2.4 Requirements Categories

The end-to-end requirements can be classified into six general categories as follows:

1. Foundation Requirements
2. Offline distribution of content and content descriptions
3. Locating content by Content Users
4. Negotiation of content delivery from source to target
5. Delivery of content from source to target
6. Execution of post delivery management

Categories 2-6 are actually sequential phases of an end-to-end “experience” where a Content User receives a menu of content to choose from, picks a particular piece of content to consume and then receives it.

These requirements will be submitted to the appropriate standards bodies (e.g. IETF, W3C, etc.) for consideration in organizing the necessary technical protocols and other technologies to support broadband content delivery. The requirements will also be submitted to the content creator and content provider communities for their evaluation and endorsement.

In addition, the business group within the BCD Forum is identifying a set of business models for the flow of information. Our goal is to test the requirements we have identified for the delivery models to see how well this supports the business models.

Following this above structure, this document is divided into seven sections, the first section describing the vocabulary used in the document, and the rest discussing the requirements for each category from the aspects of the four functional groups.

## **2.5 Vocabulary**

*Content:* A conceptual identifier of a media object and its constituent digital bit streams. The identifier is a movie title, a song name and in the case of a live event, an event identifier. Note that content could be comprised of several different bit-streams. For example, a particular movie could be encoded in three different qualities and have several language tracks. A live event could also be encoded in different qualities.

*Authentication Server:* The server that is responsible for authenticating the content user.

*Authorization Server:* Determines if an authenticated user is authorized to view content with a particular set of delivery parameters. May also bill the user for use and generate a billing record that is sent to the content provider. The content provider or a third-party could own this server.

*Content Server:* The server that is responsible for sending the stream to a target device over some network. This may be the origin server or a surrogate server owned by the content network.

*Content Distribution/Delivery Network (CDN):* A collection of network elements arranged for more effective delivery of content to clients. Typically a CDN consists of a request-routing system, surrogates (a content server other than the origin server), a distribution network, and an accounting system. (Definition from IETF Content Peering group)<sup>4</sup>

*Edge Server:* A network element close to the edge of a backbone network that is able to perform content services. Service examples include caching, video/audio stream serving and splitting, ad insertion, transcoding, content listing, etc.<sup>5</sup> A proposed Edge Server technology can be found in two OPES documents.<sup>[6 7]</sup>

*Peering:* A collection of network elements supporting some form of interconnected operation among two or more entities owned by separate organizations. Examples include accounting peering, content list peering and distribution peering.<sup>1</sup>

*Request Manager:* The system that responds to the users request to view a particular piece of content. This is usually determined by the URI of the content. It could also be a fixed proxy. This system could be owned by the content provider or content network.

### **3 Foundation Requirements**

The following are general requirements needed to support an enhanced “experience”. Note that several of these requirements involve setting standards for describing different entities, such as content, users, capabilities, etc. These standards must be created to support existing environments and be extensible to future environments.

#### *There must be a standard way to describe content*

Metadata includes attributes that are related to the specific content, such as title, author, size, classification etc. This information enables service providers to accept, reject, personalize, filter, and classify content. The Metadata will be included for various types of content (video files, games, etc.). Data elements will enable control, blocking, classifying, finding, selecting, peering, rating and negotiating of content delivery functions.<sup>8</sup> Metadata must be able to pertain to multiple content objects so similar objects can be “grouped” (i.e. priced, rated, etc.) together.

#### *There must be a standard way to describe the demographics/preferences of users*

So that service providers are able to personalize content, users need to be able to transmit profile information about them. Characteristics such as demographic information, rating information (i.e. child, teen, adult), preferences about quality of service, and accounting/billing information could be included in the profile.

#### *There must be a standard way to describe the capabilities of devices and capabilities of network connections*

Device profile information and network connection information is needed so that content and service providers can optimize the format and delivery of content for a specific device and network. For example, a device may have a specific screen size, which dictates the maximum size of an image that it can be viewed, whereas the network capability may dictate the maximum bandwidth video content that can be viewed in real time.

#### *There must be a standard format for describing accounting events*

Appropriate accounting information must be logged during all phases of broadband delivery. This includes summary reports from subsystems supporting delivery profile subcontracts. This accounting information can be used for both statistical purposes and for billing purposes. Examples include the content provider billing the Content User for viewing content, or the network provider billing the content network for network

## Broadband Content Delivery Forum

services. Since every organization in the value chain needs to recognize revenue, the particular accounting information that they need must be made available to them.

This format must include a set of identifiers that allows related records to be identified and grouped together based on some criteria. This creates an “audit trail” for a particular end-to-end “experience”.

In addition the BCD Forum has identified the following business issue related to auditing and accounting:

- Who is accountable to the end user to make it all happen – e.g. manages the business rules, ensures delivery, and coordinates the various participants along the delivery path. One rule of thumb is that the organization with final accountability is “the guy who collects the money.” This key participant is the customer for audit trails.

### *There should be a standard way for distributing accounting information*

Standards are needed to ensure the availability of accounting information and security of information so it is not compromised.

### *There must be a standard protocol for authentication of data senders and users*

This must be done before, during and after content delivery. Only trusted organizations should be able to examine user’s private information. Any private information must be sent using secure transaction systems. Content providers should be able to authenticate the systems (intermediaries and Content Users) to where they send data.

### *There must be a AAA infrastructure*

In this document, authorization, authentication, and accounting requirements are kept separate and assumed to be performed by separate functional entities. These functions, however, are intimately tied together. In reality, there must be “glue”, a AAA infrastructure that binds the different functions into a cohesive system.

### *Content and services should be able to be personalized*

Based on the user’s profile or device profile or some other inferred information, content should be able to be customized. This could include customization based on geographic information, etc.

### *There must be a standard way for managing and delivering distribution rights*

There are two distinct models for rights management when data is redistributed:

- 1) Trusted agents (e.g. content networks) take on the role of retailers or value-added resellers, managing usage rights (within specified limits) and providing related accounting information up the value chain.
- 2) Digital rights management is implemented end to end, by separating the distribution of content from the distribution of the ability to use the content (e.g. encryption keys) and conducting the “ability to use” distribution as a direct transaction between content provider and end user.

*There should be a way to create a sequence of separate events that create the "experience"*

It should be possible to chain the delivery of content. After one session is completed, the content provider (or even the content network ) may want to offer other value-added services such as recommending other content which lies in the same category as that just viewed, e.g. if the user has just watched a spy film, other spy films could be recommended. The content provider may consider this an incidental follow up to the main content, but the network provider sees this as more billable traffic. A content network may see this as two separate billing events, one to the primary content provider and the second to the trailing provider.

## **4 Offline Distribution of Content and Descriptions**

Since original content creators are normally not the entity that serves content to Content Users, content and descriptions of content are first distributed to the appropriate entities that ultimately service the user. This could be either content providers with content servers or content network.

### **4.1 Content Provider**

There are several requirements of the content provider needed to provide the information needed for the user to ultimately to see the list of available contents; this may be done by the content server or edge services. The list of available content is not constrained to be pre-recorded media; it can also be media generated during live events.

*The names of the available content needs to be generated and distributed*

A "name" is simply an identifier that describes how to find the content and/or its associated Metadata. This is probably a URI (universal resource identifier). Included within the name could be user requirements (such as cost) that the user is expected to decide upon if the content is desired. This may or may not be part of the metadata for the content.

*The Metadata needed to access the content needs to be made available along with the name*

This could include quality of service (QoS) requirements, payment rules, content classification, and security issues. For example, one payment rule could require that users pay for content before the delivery is started. Another rule might specify that payment and the starting of the content occur in parallel. In addition there may be different "treatments" to the content, such as watermarks, DRM special actions, etc. The content provider will have done these and specified this in the metadata. The metadata will also contain information about the *location* of the data. This could be a content provider identifier, an address for receiving a multicast event, etc.

*The Content Provider should have the ability to control the modification, access, deletion, and distribution of content to the Content Network in a secure way*

The content provider may wish to invalidate, modify, or distribute cached copies of both content and metadata and use a secure channel. Distribution requirements include the ability to automatically schedule the replication of content to specific destinations.

## **4.2 Content Network**

The Content Network should be able to pre-load content and/or content metadata to its content servers

There must be a method for allowing the content network to request content and metadata from the content provider or for the content provider to pre-distribute content to the content network.

## **4.3 Content User**

The Content User is not directly involved during the pre-distribution of content.

## **5 Locating Content**

The requirements in this section deal with a Content User who wishes to view a menu or list of available content for a particular content provider or content network. In addition, the requirements are the same for either the content provider or content network since both can “serve” content to the Content User.

### **5.1 Content Provider/Content Network**

The user needs to be able to find and view the list of contents

The user needs to see the list of contents, potentially in different formats. Either a content provider or a content network could generate the content list. (Note that the host of the content list does not necessarily host the actual content; these may reside on different servers owned by different organizations). The function to be executed is to determine a list of names of contents and presenting them to the user. We noted that a Web page is one standard means of this presentation. Other views include voice menus and TV screens. Regardless of the presentation medium, the potential list of contents is likely to be long; so some criteria can be employed in organizing and filtering this list on a per user basis. The filtering could be based on content metadata like publication date or price, and customer preferences or customer device capabilities, etc. For example, it may not make sense to list a high-bandwidth 1Mbps video to a user with a 128Kbps ISDN line or to list fee-based content to a user who only views free content.

### **5.2 Network Provider**

There must be a network lookup service

The content list specifies the available content. There must be a mapping from the name of the content specified in the list (probably a URI/Universal Resource Identifier) to the content provider or content network that can serve the content to the Content User. For example, the Domain Name Server (DNS) protocol provides this functionality.

### **5.3 Content User**

A profile should exist that describes the user, his/her associated devices and access capabilities

The user could provide some kind of identifier (login, cookie, etc.) that leads to a profile that the system can use to determine user information. The information may be stored at the client's device or may be held by a surrogate in the network. A user's profile may be comprised of multiple parts, information about the person (i.e. demographics, location, billing preferences and relationships), information about the user's device (i.e. screen size, bandwidth, etc.), and information owned by a service provider (history of access, payments, etc.) After that, the user wishes to obtain a list of available contents (possibly customized to the user's particular interests) from which the available content can be found. This can be presented in one of many forms including web pages, TV, or even a voice menu. Information about the user's wishes can be used to organize the content listings.

*Searching tools should exist to help the user find content*

There is a potential "long list" of available content; consequently, there is a need of some searching tools and information organization to simplify the searching process.

## **6 Negotiating the Delivery of Content**

This section deals with the delivery negotiation phase of the broadband "experience. At this point it is assumed that the Content User has selected the content to be delivered. It is possible that the user, when selecting the content, also selected explicit parameters dealing with delivery, such as quality of service, bit rate, etc. Thus, parameters needed for the delivery that were not chosen explicitly in during the location phase, are generated in this phase.

There are four main sub-phases that occur during the *negotiation* phase.

1. The user's request is routed to a *Request Manager* that handles the request. The *Request Manager* first identifies the user. The user's profile is located and the user is authenticated, either by the *Request Manager* or an *Authentication Server*.
2. Based upon the user's profile (personal and/or device profile), requests made during the first phase, user interactions and the metadata of the requested content, a delivery profile is generated by a *Request Manager*. Possible profile parameters include the quality of the asset, language preferences, and whether or not advertisements are inserted into the content. The *Request Manager* may, in fact, communicate with other entities for performing specific functions to generate the delivery profile.
3. The user may need to gain authorization and be billed for consuming the asset with the given profile. The authorization/billing may occur during either/both the *Locating Content* phase or the *negotiation* phase and could be performed by a *Request Manager* or an *Authorization Server*.
4. A *Content Server* is chosen by the *Request Manager*, which in turn, negotiates server and network resources to deliver the content from source to target.

The *Request Manager* may not be owned by either the original Content Provider or the Content Network. A *Virtual Content Provider* could own it where it simply acts as a middle man and negotiates with all parties to setup the delivery, billing, authorization, etc. of the content.

We now look at the negotiation in more detail from the perspectives of the four groups involved in the “experience.”

## **6.1 Content Provider/Content Network**

Both the Content Provider and Content Network can provide the same set of services during the negotiation depending on which type of organization is responsible for actually delivering the content to the Content User. In some business models, the content provider owns the content servers, but in other business models, the content network owns the content servers.

### *There should be a way to authorize or deny a user's request*

The content provider may use a third party *Authorization Server* to make the decision. If the request is authorized it can possibly be billed. The billing could be done by another third-party server and based on a subscription basis, per use, per bits sent, etc. Thus, the content provider must have the ability to identify the form of payment. The content provider could also outsource this functionality to the content network.

### *A delivery profile should be generated*

Once the request is authorized, delivery requirements are negotiated and a delivery profile is generated. The profile is generated from a matching of the user's profile (network connection, viewing device, etc) and the metadata describing the content. The profile contains pertinent information that is to be negotiated, such as bit rate, frame rate, etc. In creating the delivery profile the CP/CN may need to negotiate with the network provider to be charged for the delivery of the content. There is probably some service level agreement between the CP/CN and the network provider but this could possibly be done on an ad hoc basis.

### *It should be possible to specify value-added services that are executed during delivery of the content*

Other negotiations may also take place between parties. For example, the content provider may want local advertisements inserted into the content for some customers and not others. Or it may ask the content network to provide security “treatments” to be inserted into the content. This is a negotiation between the *Request Manager* and the content provider and possibly based on the profile of the Content User. Or the content provider may specify that no services may be performed on the content. These services could be based on user's preferences and authorizations. The content provider must be able to completely dictate what can and cannot be done to the content by intermediary systems.

### *Content server(s) must be chosen*

Once the content is chosen, the user is authorized and the delivery profile is generated, the *Content Server(s)* must be chosen. This could be a server owned by the content provider or a server owned by the content network or a server owned by a peering content network. Finding the content server could be done by DNS or a higher layer intelligent routing system. Note that there should be no restrictions on the number of content

servers used. Content could be assembled from components sent from disparate servers. For example, movie content could be streamed from one server while advertisement content is streamed from a different server. Also, knowledge of the actual *Content Server* may or may not be hidden from the user.

The *Content Servers* are queried to determine if it can satisfy the request. A directory of *Content Servers* needs to be stored somewhere in the network. The chosen *Content Servers* must have sufficient server resources and be able to negotiate and provision a network path to the client. If the servers do not have the content, they must negotiate delivery from other *Content Servers* that may hold the asset. The content could be streamed or downloaded from one server to another.

If the *Content Server* cannot satisfy the request, other possible *Content Servers* may be queried. If no adequate *Content Server* is found, the request could be denied or the negotiation phase could be restarted with a modified delivery profile. Once a valid *Content Server* is found to satisfy the request, the server firsts needs to allocate its own resources, such as CPU, Memory, and I/O.

*Accounting records should be generated*

Upon completion of this phase, messages need to be generated and sent to the content provider or accounting aggregator to confirm the delivery profile agreed to, against which measures or audits may be performed.

## **6.2 Content Network/Network Provider**

*A path should be provisioned between content server and Content User*

If the network allows it, network resources are provisioned on elements from source to target based on the delivery profile. Paths can be provisioned between source and target at both the content level within a content network (tunneling, application level multicast) or packet level (RSVP, diffserv, etc.). In a best-effort network, no provisioning may take place. Note there may be the notion of a “path coordinator” which establishes the path (This is discussed in another BCDF document, “*Requirements for the End to End Management of Network Quality of Service*”<sup>9</sup>). For example, the user could negotiate QoS parameters from both the content network and network provider. Provisioning could involve QoS features and the following issues:

- Finding a path or set of paths and characterizing the facilities along the way, including the Content User’s network connection and display device. Different paths may incur different costs. For example, a path with guaranteed QoS may be more expensive than a best-effort path. The users profile or the user him/herself may need to be interrogated to choose amongst multiple paths.
- An end-to-end path may include multicast capabilities at the hardware level (IP multicast) or the application level (content network overlay multicast).
- A key issue is how an end-to-end constrained path is established in a network of networks both between network providers and content networks. Peering agreements are required in this case.

- Network provider may establish pre-provisioned paths and characterizations that they are authorized to use on demand.
- A secure link may need to be established for the delivery.
- Is there a need for a network or service provider “owner” of the end-to-end path?

Once the *Content Server* is chosen and the network is provisioned the next phase of the broadband “experience” begins, the *Delivery* phase.

### **6.3 Network Provider**

*There should be a standardized protocol for requesting guaranteed or prioritized bandwidth from a network provider*

This will allow a content provider or “path coordinator” to establish an end-to-end delivery path from different network providers. If a network provider has the capability for providing QoS paths, it should implement this protocol.

### **6.4 Content User**

*A profile should exist that describes the user, his/her associated devices, and access capabilities*

The user’s profile is used in several different phases of the broadband “experience.” At the start of the negotiation phase, the *Request Manager* requests user profile information. The metadata for the content may contain attributes of service delivery (QoS), device/software requirements, and attributes on how to handle billing and payment. The user profile and content metadata are matched and the delivery profile is generated.

*The user should be able to interact with an agent to determine the best method for delivery*

The user may need to interact with the *Request Manager* or *Content Server* (or some other entity) to specify delivery parameters. For example, a user may wish to pay more to view an asset with no advertisements.

*There should a method so the user can be billed*

A major part of this interaction between the *Authorization Server* and the Content User could be the billing aspect. This interaction could be performed manually between the server and actual user (authentication must be performed first) or done by a communication transaction between the server and a user’s agent that represents the user. Billing requirements may include the following:

- A digital ticket may be exchanged between the user and the *Authorization Server*.
- The content or service provider may require prepayment, or special logging.
- The user may need to specify a password and/or give billing information to gain access.
- A transaction record may be generated with authorization/billing information
- Billing information sent between the *Authorization Server* and Content User must be secure.

## **7 Delivery of Content from Source to Target**

This section outlines the next phase of the content delivery process. In the previous phases, the content has been found and the delivery profile, which includes path discovery and possible QoS negotiation, has been determined. These requirements deal with the delivery of the content and the monitoring of the content delivery and any required sub-sessions.

### **7.1 Content Provider**

The content provider should have a way to be contacted in case of delivery changes

For the Content Provider (which is the content creator or its delegate), its role is very much passive in this phase if a content network is responsible for delivering the data. The asset is found and the delivery profile negotiated prior to this phase. However, the content provider may become involved as the result of a fault or change order message delivered by the content network. To this end, the content provider may be listening for messages during this phase (as determined by the service level agreement between content provider and content network). If the content provider is the owner of the content server, it will be listening for messages from either the user or any other party involved in the delivery of the content.

### **7.2 Content Network**

The content network should be able to perform manipulation or customization based on direction from the content provider

If, for any reason, enhancement, customization, or manipulation of the content delivery is needed (for example, for trick play, local ad insertion, video transcoding), the content network will listen for messages and will allocate any additional resources. For example, in the case of local ad insertion, the main program will probably originate from a different server than the local ad content. The content network is responsible for knowing when to trigger the local ad insertion sub-session as well as for the allocation of resources needed to deliver the ad. The content network may also need to collect data on which ads are delivered, to whom they are delivered, as well as at what times, for later auditing purposes. In all cases, the content network must adhere to the policies of the content provider for performing customization or manipulation of content.

The content network should have the ability to realize monetary value from the Content Provider

Content networks will provide the delivery and management (transaction, refund, settlement, collections) services for the content and they need to be remunerated. This will entail keeping accurate accounting records for input into a billing system. These accounting records must be able to be delivered both in real-time time and incrementally in an offline batch mode.

Content should be able to be distributed/streamed across independent content networks

The path between a content source and target may actually be a chain of independent, yet associated content networks. Independent in the sense that each controls the resources within its domain; associated in the sense that some peering arrangement that allows for negotiation of resource allocation has been previously established. Taking such a case, the need for standard methods for negotiation and monitoring of content delivery events is even more imperative.

### **7.3 Network Provider**

The Network Provider has the major responsibilities during this phase. After the delivery path has been determined, the appropriate resources allocated, and security policies have been set up, the NP starts the delivery of the asset according to the negotiated delivery profile.

*The network provider should be able to be contacted in case of delivery problems or changes*

The NP will also listen for messages that may indicate a fault or change in the content delivery. Fault messages can include anything from path overload up to a total failure of the content delivery. Messages could also indicate that user's have been removed from a multicast group. The NP will receive such messages, analyze them and take appropriate action. In the case of a failure of the content delivery, the NP may simply attempt to wait during a timeout, and attempt to restart the content delivery. If the fault is determined to be severe enough, the NP may also redirect the content delivery, either changing the serving location and/or changing the path (subject to renegotiation). Or the NP may terminate the content delivery and refuse to re-establish, if no replacement resources can be allocated. In this last case, some kind of refund procedure may be followed during the auditing phase.

The NP will also listen for change orders from other involved parties. Such orders can originate from any entity in the value chain. In the case of a CP, the change order may occur, for instance, when content is to be delivered from different servers. Change orders may also originate from the Content User, such as in the case of trick play. In this case, the NP will pass along messages to the appropriate content server that will provide VCR-like functionality, such as fast-forward, rewind, or pause. These messages may, in fact, not go to an originating server, but perhaps to a cache server, or even a specialized index server, which handles trick play within a content network, in a real time sub-session. Also, the Content User may need additional authorization in order to receive the additional service. Whatever the case, the NP will determine what resources are necessary and allocate the resources at the proper time.

*The Network Providers should have the ability to realize monetary value from the Content Provider and/or the Content User*

Network providers will provide the delivery and management (transaction, refund, settlement, collections) services for the content at the network level and they need to be remunerated. This will probably entail keeping accurate accounting records for input into a billing system.

*Content should be able to be distributed/streamed across independent network providers*

It should be noted that the above perspectives take the simple case of the network provider having end-to-end control of the resources. However, the network provider may actually be a chain of independent, yet associated, core and access providers in the delivery path. Independent in the sense that each controls the resources within its domain; associated in the sense that some peering arrangement that allows for negotiation of resource allocation has been previously established. Taking such a case, the need for standard methods for negotiation and monitoring of content delivery events is even more imperative.

## **7.4 Content User**

*The user should be able to interact with the service*

The Content User's experience during this phase can be seen primarily in the terms of customer satisfaction. Is the content being delivered in a satisfactory manner? If not, the user may issue a change order. The user may ask for better service, perhaps even explicitly ask for a greater allocation of resources. Or conversely, the user may ask for lower-quality content to be delivered over the same path.

Also, the Content User may request changes that manipulate or enhance the content, such as in trick play. Or the user may be interacting with content, such as in the case of an MPEG-4, when a "clickable" ad object may be placed within the frame itself. When the user clicks on the object, the content delivery may be redirected, with content being delivered from a new location.

*The rendering device should generate accounting assets*

Also, client software may monitor the content delivery and determine if the delivery profile requirements are being met, and if not, take appropriate action. This client software may also log content delivery events and pass such data to other entities in the value chain, for use in the auditing phase.

*Notification should be generated if the user expectation falls below the negotiated profile*

During delivery some negotiated QoS parameter could be violated. For example, average bandwidth delivered might fall below some threshold. The content provider and/or the content network should be notified in this case so that the path could be changed, or a refund given.

## **8 Execution of Post-delivery Management**

This section describes the post-delivery processing that is carried out following the reception of an item of content. It describes the last phase of the "experience".

### ***Definition of End of Session***

The end of the session can be determined either by the natural end of the content or by the forced termination by any of the parties involved.

The Content User could terminate the content because he/she is no longer interested in continuing. This could be because after a free viewing period, he does not wish to pay for the full content, or simply because he does not wish to continue regardless of any payment already committed.

The Network Provider or Content Provider could terminate the session due to expiration of credit, equipment or network failure.

## **8.1 Content Provider**

### *The content provider should receive accounting information*

The Content Provider may have had no involvement so far, since the content network may have delivered the content from a surrogate server without its knowledge. Thus, the Content Provider now needs to be informed of what happened during the session so records can be updated. This information should include a report about the network delivery and a report about the user's activity. The information could also be sent to an aggregator who summarizes the information in a report format. If the Content Provider is not the Content Creator, other accounting records may need to be generated.

Whether or not the Content Provider is billing the user directly, accounting information may still be important for statistical purposes. The information will include all of the events that occurred during the session, including information like start time, overall viewing time, whether all of the content was viewed, etc. Both the content network and network provider will deliver this and other pertinent information. It is possible, however, for the content network to aggregate the information from its network providers. Again, the appropriate accounting records must be correlated in some way to identify an end-to-end transaction for creating an audit trail and performing reconciliation.

### *The content provider should be able to bill the user either directly or indirectly*

If the Content Provider is billing the user, some kind of user handle will be needed to build profiles of viewers based on their choice of content. In this case, the identification does not have to include an actual username – it simply has to uniquely identify a particular user for correlation purposes. This information can then be used to provide the kind of services described in the Content User section above, allowing targeted content recommendations to be made in the future. We emphasize that there are cases where the actual identity of the user does not have to be known by the content provider.

## **8.2 Network Provider**

### *The network provider must close the session and free any resource*

The Network Provider is responsible for any sessions that were established for the purpose of delivering the content. These will now be closed and any allocated resources will be freed, ready for re-use.

Accounting records should be generated and sent to the appropriate accounting manager

As the session has progressed, the network provider will have built up accounting records, which will be complete once the session has terminated. This can now be processed and sent to the content network and/or content provider. Security measures must be in place for the transfer of this information.

### **8.3 Content Network**

Accounting records should be generated and sent to the appropriate accounting manager

The content network may provide information to the content provider of which surrogate (if any) actually served the content, which mechanism was used to deliver the content (e.g. tunneling) which encoding was used (in case of multiple encodings being provided), and statistics about the session, etc. This information must be correlated with other accounting events that were generated for the same session. Finally, the content network will either redirect the user to the content provider to present billing information, questionnaire etc, or provide the function itself and forward the results to the content provider.

### **8.4 Content User**

The User's rendering software should make available any statistical or informative records that were generated

When the session ends, the client player should send accounting records to the content provider or an accounting aggregator.

The user could be offered a questionnaire

After having viewed the content, the user may be offered a questionnaire by any of the entities in the value chain, asking for feedback about what has been seen. The complexity of this questionnaire is a balance between getting a lot of information and minimizing the disruptive effect on the user. It could just be a score out of 10. A financial incentive could be offered for filling in the questionnaire. The questionnaire could be different depending on whether the user terminated the content early, or viewed it until the end. The user's profile could include rules about when questionnaires are offered.

A summarization of the session must eventually be available to the user

The user must be presented a billing summary for the content just viewed, breaking down the charges if the model is not flat rate. This report may not be immediate; it could be an aggregation of sessions that is sent periodically, i.e. a telephone bill.

If the content delivery was terminated prematurely by any party other than the user, then some kind of explanation should be given. Obviously, applicability of billing information and questionnaires should be carefully considered in this case.

**Note**

In summary, the post-delivery processes depend very much on the accounting and billing models as well as the services that the various parties wish to offer. This and also the demarcation of tasks between Content Provider and Network Provider(s) can complicate the issue. However, the main components will usually be constant. These are:

- Graceful session teardown and release of allocated resources
- Processing and delivery of billing information
- Receiving user feedback
- Gathering and correlation of statistics

## **9 Authors Addresses**

Michael Vernick (Primary Contact and Working Group Co-Chair)  
Bell Labs/Lucent Technologies  
[vernick@bell-labs.com](mailto:vernick@bell-labs.com)

Simon Bryden  
Nortel Networks  
[sbryden@nortelnetworks.com](mailto:sbryden@nortelnetworks.com)

Michael Condry  
Intel Corporation  
[condry@intel.com](mailto:condry@intel.com)

Gary Disher  
[g.disher@snet.net](mailto:g.disher@snet.net)

JC Straley  
Sprint Labs/Sprint Corporation  
[jcs@sprintlabs.com](mailto:jcs@sprintlabs.com)

Wil Walkoe (Working Group Chair)  
Sprint Corporation  
[Wil.walkoe@mail.sprint.com](mailto:Wil.walkoe@mail.sprint.com)

## **10 Acknowledgements**

We would like to acknowledge the following people for their insightful comments in producing the document: Elon Cosla (Optibase), Thomas Nojd (Telia), Per Tholin (Telia), Eitan Tobi (Wydeband), William Ho (Reliacast), Solom Heddaya (Infolibria), Abbie Barbir (Nortel), Gene Chang (Funk Software), and Jay Goldin (Digital Fountain).

## **11 References**

<sup>1</sup> M. Day, B. Cain, G. Tomlinson, “A Model for CDN Peering,” draft-day-cdn-model-04.txt (work in progress), January 2001, [URL:http://search.ietf.org/internet-drafts/draft-day-cdn-model-04.txt](http://search.ietf.org/internet-drafts/draft-day-cdn-model-04.txt)

<sup>2</sup> Siegal, Steven, “Network Supported Applications,” Bell Labs Technical Journal, Vol. 6, No. 2, July/Dec 2001.

<sup>3</sup> Recommendation ITU-R M.1079, “Performance and quality of service requirements for International Mobile Telecommunications-2000 (IMT-2000).”

<sup>4</sup> M. Day, B. Cain, G. Tomlinson, “A Model for CDN Peering,” draft-day-cdn-model-04.txt (work in progress), January 2001, [URL:http://search.ietf.org/internet-drafts/draft-day-cdn-model-04.txt](http://search.ietf.org/internet-drafts/draft-day-cdn-model-04.txt)

<sup>5</sup> Beck, et. al., “Example Services for Network Edge Proxies,” draft-beck-opes-esfnep-01.txt, November 2000, URL <http://search.ietf.org/internet-drafts/draft-beck-opes-esfnep-01.txt>

<sup>6</sup> Open Pluggable Extensible Services Web page, <http://www.ietf-opes.org>

<sup>7</sup> Tomlinson, et. al., “Extensible Proxy Services Framework,” draft-tomlinson-epsfw-00.txt (work in progress), July 2000, URL <http://search.ietf.org/internet-drafts/draft-tomlinson-epsfw-00.txt>

<sup>8</sup> E. Goshen, “Metadata Requirements for End-to-End Content Delivery,” BCD Forum Contribution, June 2001.

<sup>9</sup> S. Bryden, G. Disher, M. Fisher, D. Teichman, W. Wilkoe, S. West, “Requirements for the End to End Management of Network Quality of Service,” BCD Forum Contribution, February 2001.