

Content Reference Forum

Introduction



Title: Content Reference Forum Introduction
Document: CRF-004
Version: 2
Date: March 17, 2003
Authors: Content Reference Forum
Status: CRF Document

1.0 THE PROBLEM

Viral Marketing & Distribution

Consumer to consumer (viral) marketing & distribution is already a reality. Content is moving around electronically. Unfortunately, it is unreliable and ungoverned. Consumers send URLs to each other and trade content over various P2P networks. URLs that point to specific assets are often broken. Files acquired over P2P networks are inconsistent. Digital content related markets (content, hardware, software, and services) would benefit greatly from reliable, predictable, easy to use consumer to consumer marketing and distribution.

- ➔ How can we support viral marketing & distribution without forcing consumers to exchange large files that may or may not work when they arrive?

Appropriate Compensation

In today's viral marketplace, authors are not compensated and their rights are not protected. Additionally, there are others whose rights are not accounted for. Retailers, promotional partners, fulfillment services and consumers are among those who may be owed compensation -- financial or otherwise -- based on contractual relationships (including click licenses). Also, new types of relationships could be defined to create incentives for participation in the marketplace if only there was some way to track value chain participation¹. How do we stimulate consumer sharing as a promotional activity while still protecting the rights of authors and interests of others in the value chain? More specifically in this context:

- ➔ How do we track meta-data so as to allow the viral marketing and distribution system to track and protect value chain interests?

Digital Distribution and the Global Economy

If we consider the digital distribution of content, no single file will be appropriate for distribution to all users. Many content elements (e.g. songs or movies) have different rights or controls associated with them in different territories or for different uses. Content may not be available to all consumers in the same way due to license or territorial restrictions, or even distribution channel availability. Some content elements will be in different formats depending on the intended use or means of delivery. Streaming a 9-Megabit DVD audio stream to a listener might be appropriate using Internet 2 but would not be appropriate over GPRS. Material might be available using different codecs. The rendering device or user preference should determine the correct choice. If a consumer comes to a web site, they should not have to know what formats are appropriate to their local configuration; instead, this should be handled automatically and invisibly. A service provider should not worry about creating custom interfaces to each content provider or retailer, and vice versa; instead, they should be able to leverage a set of interoperable e-business standards.

- ➔ How can we make content links and content distribution effective regardless of the medium, the channel, the platform or the renderer's location?

¹ Value chain participants are any entities that take part in the creation, distribution, sale, or promotion of content, as well as the consumers who virally market/distribute the content and the consumers who acquire the content.

- ➔ How can we leverage existing and create new e-business standards and apply them to content distribution while ensuring interoperability for developers and system's participants?

2.0 THE REQUIREMENTS

Overview

To summarize the points above, we are trying to create an environment that enables the following:

- ➔ Efficient support of viral marketing and distribution in a manner that provides results tailored to individual users based on their particular location, platform, preferences, etc. and based on terms determined by the affected value chain participants
- ➔ Support of new business models that take the value chain participant interests into consideration
- ➔ Leverage common electronic business and digital content management standards for end-to-end content distribution networks

Therefore, we require:

1. A means of referring to content in the abstract, rather than to a specific file or instance of the content. We will call this a "Content Reference".
2. A means of tracking value chain participation
3. A means of tracking value chain relationships and interests therein (e.g. contracts) that should be respected in transactions resulting from use of a Content Reference.
4. A means of obtaining information about the environment of a consumer who is trying to exercise a content reference.
5. A means of resolving Content References given: a Content Reference, value chain information, and context (consumer environment information).
6. A means of expressing resolved Content References such that consumers can locate, transact for, and acquire instances of the referenced content.
7. An ability to select implementation profiles such that developers of Content Reference architecture-based services can be assured of interoperability with other services

3.0 THE SCOPE AND VISION OF THE CR FORUM

The goal of the CR Forum is to facilitate the distribution of legitimate content in peered environments with the least friction possible. We want to dynamically deliver the right content with the right commercial terms for a given consumer and value chain context. Therefore, we need a way to refer to content in the abstract, a way to describe user and value chain contexts, and a way to define and process the rules that will determine the appropriate instance of content and commercial terms to deliver.

Standards bodies and industry initiatives such as MPEG, W3C, DOI, WS-I and OASIS are working on portions of the solution and many of the pieces are being defined. Unfortunately, there are pieces missing and there are no agreed mechanisms to tie the different pieces together into an interoperable whole.

A common vision has emerged for the creation of a Content Reference Forum, which will synthesize related specifications from existing standards efforts and fill in the gaps as necessary. This initiative is founded upon the following fundamental principles:

- First. A solution to this problem will enable new business models and increase commerce around digital media.

- Second. The problem of providing the right content with the right offering is multifaceted and complex.
- Third. It is not realistic for a single company to develop detailed specifications. Success requires collaboration of the organizations and disparate industries that will share in the benefits.
- Fourth. Standardization of any related specifications will be key to the successful and widespread adoption of content reference and resolution systems.
- Fifth. The forum focuses on practical aspects of building commercial systems for content distribution compliant with the Content Reference architecture. To that effect, the forum will work to leverage pervasive infrastructure technologies and establish interoperability guidelines and specifications for systems' deployment.

The intent of the Content Reference Forum is not to supersede other standards activities but to integrate their results into a practical and implementable architectural framework.

Forum Goals

The primary work product of the Content Reference Forum will be a set of specifications that describe how to build content references and the systems that process them. Specifically, these specifications will describe how to build systems that:

- Construct and manipulate content references, which identify and describe content in the abstract
- Describe user contexts (supported formats, rendering capability, territory, etc.)
- Identify value chain participants and describe the contractual obligations between them
- Express contractual obligations governing distribution and use of content
- Resolve abstract content references given user context, value chain obligations, and content rules and produce offers for specific content instances
- Identify, find and deliver appropriate content
- Enable processing of offers and clearance of content-related transactions
- Enforce rights and rules for acquiring and using content

Existing Standards and Activities

Because of the ongoing work in the existing fora, many of the infrastructure pieces that CRF can leverage have already been defined or are being defined elsewhere. The Content Reference Forum will leverage existing and developing standards and specifications as appropriate and liaise with the bodies that govern them.

So far, MPEG-21, DOI, MI3P, OASIS, WS-I and W3C activities have been identified as relevant:

- MPEG-21 is defining how to declare, identify and describe content (DID and DII&D), describe user context (DIA), express and enforce the rights to the content (MPEG REL, indecs2 RDD, IPMP), and more. MPEG-21's planned work on event reporting may provide standard interfaces needed for content reference systems.
- Groups such as MI3P, DOI, cIDf, etc are developing content identification initiatives for different types of media. The DOI/Handle system is creating an infrastructure for reliably pointing to digital objects based on their identifiers.
- Activities in OASIS are intended to create standards for registering and finding value chain participants and establishing collaborative agreements, reliable messaging, business transactions, web services management and security, etc.

- W3C maintains and progresses such core standards as HTTP, XML, XML-Schema, XML-Encryption, XML-Signature, URI, WSDL, SOAP, etc.
- WS-I accelerates adoption of web services by developing interoperable profiles of constrained specifications and sets of supporting tools and samples applications

The Content Reference Forum will work to select the appropriate standards and assemble them into specifications profiles assuring interoperability between compliant commerce participants.

New Elements

While the existing standards activities define content reference components, they do not describe how to use the components to build content references and iterate them as they travel through the value chain. Additionally, systems for resolving those references, processing them according to contractual obligations ("Reference Services"), will be required.

The Content Reference Forum will develop these and other missing pieces and ultimately describe them in a set of specifications. These specifications will help to "glue" together the work of relevant standards forums into a blueprint for deploying architectures for content distribution and commerce. Solving the problem of "the right content with the right offer" will help to realize the promise of the new communications and processing technologies.

The key new elements that Content Reference Forum will develop include:

- the data and message formats for expressing content references
- contract expression language (CEL)
- reference services and RS network

4.0 The Approach

The framework for addressing the high level requirements identified in section 2.0, we propose to use the following approach.

Referring to the Content and Carrying the Relevant Data

To address the first requirement, we propose the creation of a "Content Reference." This should be a small XML-based data structure that refers specifically to a content element or set of elements in the abstract. A consumer can trigger resolution of a Content Reference received from a friend by activating (e.g. clicking on or selecting) it.

To serve our needs, a Content Reference must carry an identifier unique to the referenced content. This identifier might indicate an instance of the content that can later be resolved to an identifier if no abstract identifier is available. For example, even if a Content Reference for a song contains an ISRC number that is specific to one particular mix, we can still use this to identify the song in the abstract. The Content Reference will also need to contain whatever additional data is required for resolution (e.g. the address of the resolution service, etc.). Finally, the reference should carry textual data about the content so that human recipients can identify the content that the reference points to.

Tracking Value Chain Participants & Consumer Information

To fulfill the second and fourth requirements in section 2.0 above, we also need to track the value chain participants and consumer details. Rather than transmitting this information separately, it can be included

as part of the Content Reference. Note that since this information may have financial consequences, a Content Reference may need to be protected or tamper-resisted.

Resolving the References

To fulfill our fifth requirement in section 2.0 and resolve Content References, we need a “**Reference Service**.” This service must be able to parse Content References and return the appropriate links and other data so that consumers can select and access a relevant instance of the referenced content.

The data included in a Content Reference can only reflect a subset of the information required to effectuate an appropriate resolution. A reference can include data that the user received (the content ID, the sender, etc.) and contextual data (e.g. rendering device, territory, etc.). There are other elements we will call “**eContracts**” defined by law or contract that will need to be processed dynamically at the server to determine how and under what circumstances content instances are available. The most obvious example of this is simply the list of how content is available in a given territory. For a more sophisticated example, suppose that an author has changed distributors since a Content Reference to her work was first created. Rather than failing to find the content, a Reference Service will ensure that the right distributor is used by checking its database of eContracts to determine the appropriate distribution mechanism. This can vary by territory and even if the distributor has not changed, the original reference might have been sent from one territory to another with different contractual obligations. Note that eContracts provide us with a way to record the rights and interests of value chain participants, thereby fulfilling our third requirement.

Presenting Resolved Content References

Once a Reference Service has resolved a Content Reference, it needs to return data in a form that can be acted upon by a consumer, which is our sixth requirement in section 2.0. This “**Offer**” can be similar to a Content Reference except that it must identify the location of a particular content instance (or instances) and may include the currently available commercial terms for acquiring the content (e.g. rendering restrictions, price, expiration date of the offer, etc.). If multiple options are available, the consumer can make a selection and activate the Offer of their choice, initiating the content acquisition process or if necessary, the purchase process.

Selecting Infrastructure Technologies for Content Reference Architecture

The solutions and services described above must be deployed within an interoperable framework of pervasive infrastructure standards, which is our seventh and final requirement. We must assist developers and implementers in the creation and deployment of Content Reference services. To that end, we require careful selection of specifications based on e-business and content management standards in the areas of content identification and description, transport, messaging, service description and discovery, management and protection of intellectual property, security and authentication.

Definitions

Content Reference

A Content Reference (CR) is an XML data structure or text string that identifies a piece of content in the abstract, the value chain participants with interests in the content, and the user context in which the CR is being exercised (e.g. territory, rendering capability, service provider, etc.). CRs are resolved by Reference Services into Offers.

Content References contain:

- Content description information
- Distribution channel information
- Consumer environment information

The tables below list some identifiers from each type of category that might be included in a CR.

Content Description IDs	Description
Object ID	Uniquely identifies a specific content element (e.g. song, film, game, etc.).
RS ID	The identity (i.e. the address) of the RS(s) that can be used to resolve the CR.

Distribution Channel IDs	Description
Distributor ID	The unique identifier that refers to the owner/licensor (Author) of the referenced content element(s). Typically the company that packages the content.
Channel ID	An entity that licenses content from a distributor to other channel participants or to consumers (e.g. aggregators, radio stations, MSOs, retailers, etc.).

Environment IDs	Description
Territory	Identifies the territory in which the consumer is located or governed.
Format ID	The format(s) supported by the local environment. These IDs may indicate codecs, bitrates, file formats, or even the physical media on which the content will reside.
Platform	Rendering Platform (e.g. Linux, Windows XP, etc.)

EContract

Once the CR arrives at the Reference Service, the appropriate response is generated based on the rules expressed by eContracts. The role of eContracts is to describe the business relationships (generally contractual) between or among parties participating in the value chain. These contracts are expressed in a form readable by Reference Services that indicates how content may be made available in different circumstances. As an example, an eContract might express that an author has made their content available through Distributors A and C in one territory and through Distributors B and C in all other territories.

As long as CRs are resolved according to these rules, the interests and obligations of the value chain participants can be respected.

EContracts will be expressed using a Contracts Expression Language (CEL). CEL is similar to a Rights Expression Language (e.g. REL in MPEG-21) but is expanded to include prohibitions and obligations. CEL must also be able to express law (e.g. a behavior is prohibited in France) and preference (I am allowed to distribute through two different channels but in certain cases I prefer one over the other).

Reference Service

A Reference Service (RS) takes the data contained in Content References, uses it to identify relevant eContracts, and generates Offers for the referenced content based on the instructions in the eContracts for the specific user and value chain information.

Offers

An Offer is a data structure that expresses how a consumer can acquire a particular instance of content. An Offer may or may not contain financial terms. An Offer will typically describe a content instance and

indicate where it can be obtained. The location information should be provided in a form that can be understood by the user's player or device (e.g. a URL). An Offer may also specify that some form of payment, whether in the form of money or information, is necessary to acquire the content. Either way, the Offer should contain sufficient information that users can decide whether or not to make the commitments, if any, which are necessary to obtain the actual content. Offers will likely be expressed using a Rights Expression Language.

A VIDEO EXAMPLE

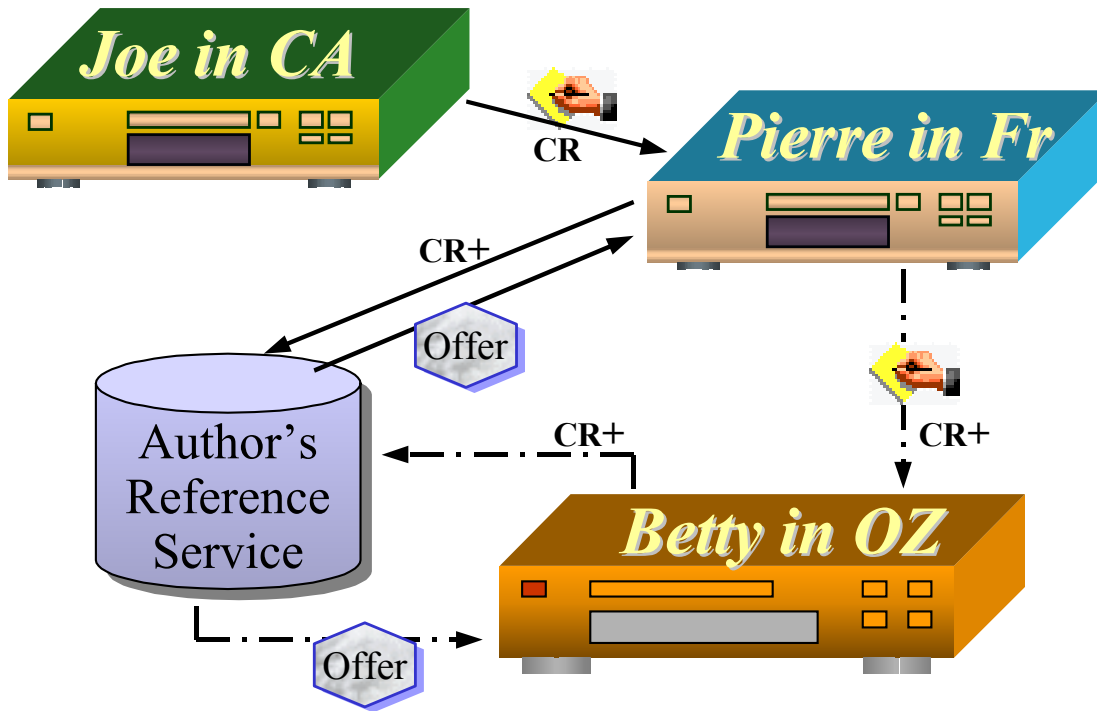
Referring a Video Program to a Friend

Imagine that Joe in California has just seen the movie “The Mummy” on HBO cable TV. He likes the movie so much that he wants to tell his friends about it. Obviously Joe is not going to send the whole movie to his friends; instead he wants to send them enough information that they can automatically find it and view it when it becomes available in their particular location.

This first thing to note is that Joe is not going to be able to send all of his friends a channel and time. He may not know when the movie will be on TV again and even if he did, it is not likely that all of his friends have access to the same channels, especially because some of them may have different services (e.g. cable vs. satellite) or live in areas where certain channels are not even available (e.g. overseas). Even if everyone did all get the same channels, the movie’s air time might be different in each location. Fortunately, Joe can solve this problem using a Content Reference.

The CR that Joe sends to his friends will contain information about the movie, such as the title, stars, and maybe a brief description, so that his friends can understand what the CR points to. It will also include enough metadata to identify the movie to the appropriate Reference Service, independent of location, language, and distribution channel. For example, the metadata might include an ISAN identifier.

When one of Joe’s friends receives and selects the CR, a Reference Service look-up will be initiated. The Reference Service takes the CR and information about the user who activated it, and returns information about the content’s local availability.



For example if Joe's friend Pierre in France clicks on the content reference, Pierre's identity, location (France) and possibly hardware & software configuration will be sent to the Reference Service along with the CR. The Reference Service can then identify the various ways that the movie is specifically available to Pierre. In this example, the reference service may determine that "The Mummy" is currently available in France on-demand from Canal+ Satellite for €4.95, or from www.amazon.fr on a region 2 DVD for €30.

Pierre can now determine whether he wants to get the content from Canal+ or Amazon.fr. Note that if Pierre likes the movie and wants to send a CR to his friend Betty in Australia, it does not matter that Betty does not have access to HBO, Canal+, or Amazon.fr, since the Reference Service will determine how the content is available in her territory.

5.0 THE WORK PROCESS

Technical work within the Content Reference Forum is done by working groups chartered by the Board. Currently four working groups have been established: Requirements and Architecture, Content References, Contract Expression Language and Reference Services. It is anticipated that in the near future an additional working group will be established to address the issues of testing, conformance and experimentation. The forum plans to conduct the first proof of concept of content reference services late this year.

Requirements and Architecture Working Group

Charter

This Working Group was tasked to 1) establish and maintain requirements to meet the goals and objectives of the Content Reference Forum, 2) define the architecture for Content Reference based systems, and 3) perform other tasks as needed in the development and deployment of Content Reference based systems.

Deliverables

This group is working on the following deliverables:

- Functional requirements for the Content Reference architecture
- Framework for the Content Reference Architecture
- Specification of the Baseline Profile for implementation of Content Reference services
- Specification of the Constrained footprint Profile for implementation of Content Reference services
- Use Case Scenarios that capture the requirements and exercise specifications for the use of Content Reference services

Current Status

The requirements, architectural framework and use case scenarios documents are in the working drafts being iterated within the working group. The draft specification for the Baseline Profile is expected this summer.

Content Reference Working Group

Charter

This Working Group was tasked to 1) establish and maintain requirements for data and messaging formats for Content References, and 2) develop and maintain specifications for those formats.

Deliverables

This group is working on the following deliverables:

- Requirements for the Content Reference specification
- Specification of the Content Reference format
- Specification of the Offer format

Current Status

The requirements document is a working draft being iterated within the working group. Proposed specification for the Content Reference format has been submitted and is being discussed within the group.

Contract Expression Language Working Group

Charter

This Working Group was tasked to 1) establish and maintain requirements for the Content Expression Language, and 2) develop and maintain specifications for the Content Expression Language.

Deliverables

This group is working on the following deliverables:

- Requirements for the Content Expression Language specification
- Specification of the Content Expression Language

Current Status

The requirements and the specification documents have been submitted and are being discussed within the group.

Reference Services Working Group

Charter

This Working Group was tasked to 1) establish and maintain requirements for the Reference Servers and Reference Services Network, and 2) develop and maintain specifications for the Reference Servers and Reference Services Network, including internal interfaces between the Network sub-components and external interfaces to other modules within the Content Reference framework.

Deliverables

This group is working on the following deliverables:

- Requirements for the Reference Servers and Reference Services Network specifications
- Specification of the Reference Servers
- Specification of the Reference Services Network

Current Status

The requirements and the specifications are being discussed within the group.

6.0 MORE INFORMATION

More information in the Content Reference Forum can be found at:

WWW.CRForum.org

Mail queries should be sent to:

info@crforum.org