Datacasting Networks:

A Service of the Nation's Public Television Stations

The mandatory transition to digital television broadcasting provides an unparalleled opportunity for a new generation of services that can improve educational attainment and enhance the safety of Americans.

Through a technology known as "datacasting," public television transmitters can do much more than distribute outstanding programming. These new digital facilities can wirelessly distribute streamed video and data files to computers and computer networks – with a capacity equal to thirteen T-1 data lines.

Because Public Television's mission is universal service, currently 99 percent of American households can receive an analog signal from one or more public television stations. This means that once the digital build-out is complete, nearly all homes, schools, government buildings and businesses will be able to receive digital noncommercial educational and public safety signals.

What Is Datacasting?

Datacasting is a one-way broadcast service, which reduces or in some cases eliminates the need for costly broadband internet connections. Datacasting allows downloading streaming video or large files to hundreds or thousands of locations simultaneously with the capacity of up to thirteen T-1 lines (19.4 megabits per second). Datacasting can provide broadband, last-mile, wireless data delivery to places that can't be reached with the Internet or where dial-up service is the norm. And when married with an existing high-speed network, datacasting can help maximize its efficiency and reduce the need for additional (and expensive) broadband.

Datacasts are encoded within the digital television signal, and then decoded by an inexpensive receiver that is easily hooked up to a personal computer, laptop, or computer network. The antenna can be as simple as a small portable antennae that sits on top of the PC (or laptop in the field), or users can receive the signal through a conventional rooftop TV antenna.

Digital datacasting is offered now, and access can be provided to all Americans within a very short timetable, because much of the digital transmission infrastructure is already in place. Each location – whether a home, school, government building or business – needs only an inexpensive datacasting receiver and antenna. This equipment is commercially available right now and is easily connected to a standard desktop or laptop computer, as seen below.

Datacasting Desktop Equipment

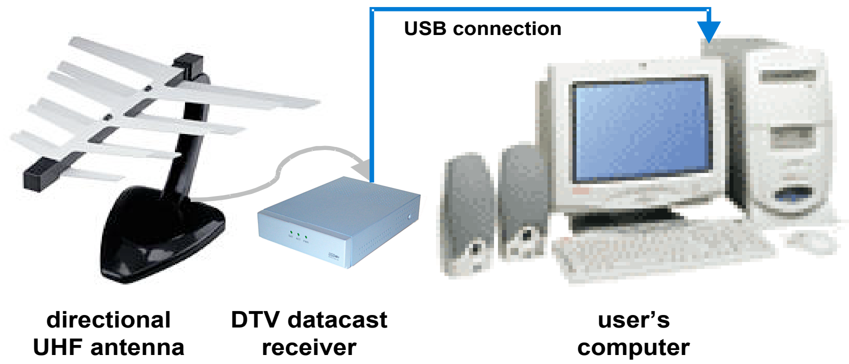

What Can Datacasting Do?

Public television digital datacasting services can provide training of first-line emergency responders, professional development for teachers and other professionals, and immediate emergency information distributed simultaneously to hundreds or thousands of sites.

The low capital investment required to receive datacasts, combined with the economics of broadcasting (where more users means less cost per user, rather than higher costs per user which is the norm in the world of the Internet) means significant economies in reaching remote locations.

Service is configured into virtual data "channels." These "channels" can be made available to anyone who has a datacast receiver (such as National Weather Service information) or can be proprietary (such as distance education courses designed to help teachers meet certification requirements). Other channels can be encrypted with the highest level of computer security (such as Homeland Security information). The bandwidth of each "channel" can easily be configured from 56 kbps to 1 mbps, or higher as needed. In a crisis situation, bandwidth can easily be increased.

The digital transition is already underway and nearly all public television digital transmitters will be on the air by May 1, 2004. The opportunity now exists to complete the build out and offer the educational and public safety benefits to all Americans.

Submitted to the Dublin Core 2003 Conference

PB Core — the Public Broadcasting Metadata Initiative: Progress Report

Alison M. White
Corporation for Public Broadcasting, USA
awhite@cpb.org

Alan Baker
Minnesota Public Radio, USA
abaker@mpr.org

Marty Bloss
National Public Radio, USA
mbloss@npr.org

Paul E. Burrows
Media Solutions, OIT, University of Utah USA
pburrows@media.utah.edu

Efthimis N. Efthimiadis†
The Information School, University of Washington, USA
efthimis@u.washington.edu

Marcia Brooks, David MacCarn, Thom Shepard
WGBH Boston, USA
{marcia_brooks, dave_maccarn, thom_shepard}@wgbh.org

Cate Twohill
Public Broadcasting Service, USA
ctwohill@pbs.org

Abstract

PB Core is the result of the public broadcasting metadata initiative (PBMI). It is an effort of the public radio and television broadcasters to develop a schema for the description of their assets. PBMI is under the auspices of the Corporation for Public Broadcasting. The paper discusses the user-centered development of the schema, the elements of the PB Core, the application profile, and the feedback and evaluation process of the schema.
Keywords: Public Broadcasting Metadata Initiative, Dublin Core, PB Core, Media Asset Description.

1. The Need for Public Broadcasting Metadata

As public broadcasting endeavors to maintain our value and values in a dramatically altered media environment, we know we must do three things: develop and deliver content across multiple platforms, strengthen our editorial and service partnerships, and engage in more efficient methods of conducting our new and legacy activities.

The recent convergence of IT capabilities with those of radio and television broadcasting has caused us and our constituents to appreciate that our prized editorial output (video clips, audio interviews, transcripts, etc.) can be understood as a series of digital assets, that can be identified, exchanged and distributed using an advanced digital infrastructure. Our ability to network – to exchange rich media content – within and across our newsrooms, production suites, satellite and terrestrial distribution systems, etc., and even with our educational and community partners (schools, libraries, museums) has never been greater. We have been afforded a tremendous opportunity for cultural relevance and operational efficiency.

In a public broadcasting system made up of hundreds of independent licensees, the challenges of organizing universal processes for asset appraisal, digitization, rights clearance, preservation, etc. are myriad, perhaps overwhelming. We did understand, however, that the foundation of any future effort in this direction would be a single, shared protocol for identifying and describing our rich media assets.

The Public Broadcasting Metadata Initiative (PBMI) is a cross-organizational, multi-disciplined effort to establish a standard for all public broadcasting content (radio and television), in order that metadata might be more easily exchanged between colleagues, software systems, institutions, community partners, individual citizens, etc. The PBMI will be a "touchstone," a single, streamlined standard to which other database structures, including those of PBS, NPR, major producing stations, and other asset/content management systems will be "mapped." It can also be used as a guide for the onset of an archival or asset management process at an individual station or institution.

The project has been extant since January of 2002, and during its first two phases of CPB Future Fund support, a team of individuals representing public broadcasting's key institutions and endeavors, along with subject matter experts (see appendix for list of participants) has worked to:

- Develop consensus regarding project objectives and timeline;

† To whom correspondence should be addressed.

- Recognize and codify the way our constituents use our content and content information. (Developed use cases based on interviews with producers, broadcast operation staff, educators, website creators, etc.);
- Examine relevant metadata standards in the media and library communities, to ascertain their applicability to our content and constituencies;
- Make information about the PBMI available via numerous conference presentations and a project website;
- Contribute and combine the substantial metadata work already performed at key institutions in public broadcasting (PBS, NPR, WGBH, KUED, MPR);
- Form a preliminary consensus regarding a single set of metadata protocols - the Public Broadcasting Core (PB Core) Metadata, Preliminary Version.1.

2. What Alternatives Were Available

The main goal of the PBMI is to create a schema that is easily understood, implemented and adopted by the Public Broadcasting community at large. PBMI embarked in a detailed review of existing metadata standards that are used for the description of rich media assets. In general, while many of the metadata standards discussed below are in development, the Dublin Core Element Set has remained stable since its 1.1 revision in 1999 [1]. Additions and other changes to the Dublin Core model come in the form of recommendations and application profiles, but the basic core of 15 elements remain unchanged. So we have built our model upon the Dublin Core that provides a solid foundation that is extensible, scalable, and easy to understand.

The standards that were considered were OAIS, SMEF-DM, MARC, METS and MPEG-7, as well as the educational standards SCORM, LOMS, IMS. These are briefly discussed below.

OAIS: Reference Model for an Open Archival Information System [2] is a framework and reference architecture for digital preservation.

SMEF-DM: Standard Media Exchange Framework - Data Model [3] is an end to end broadcast production model, workflow oriented. Our assets may involve domains or materials not exclusive or even related to broadcasting, such as CD-ROM, DVD, books. It was decided to focus on assets within a repository. Metadata was determined to describe assets as objects or files. However, SMEF mandates a specific workflow with limited options. For example, assumptions are made on the order of activities. Our experience is that productions have many different workflows that must be accommodated.

MPEG-7: “Multimedia Content Description Interface” is a highly structured standard focusing on multimedia. Our model does not preclude a station adopting MPEG-7 because the PB Core is based on the Dublin Core model and will map to MPEG-7. On the other hand, MPEG-7 is narrowly focused on multimedia, not on the wide range of other media or materials that will be found in a producing station's repository. See work of, e.g., Jane Hunter [5,6], Grace Agnew [7].

MARC: The MARC formats are standards for the representation and communication of bibliographic and related information in machine-readable form [8]. MARC requires a cataloging skill set that is not likely to be found in most public broadcasting stations. The idea of "works” focuses on the various formats or manifestations of a work that might fall under a single record. Our model insists on the integrity of each asset (version or format of the content). Dublin Core crosswalk maps to key fields in MARC http://www.loc.gov/marc/dccross.html.

METS: Metadata Encoding and Transmission Standard [9]. The METS schema is a standard for encoding descriptive, administrative, and structural metadata regarding objects within a digital library.

SCORM: The Sharable Content Object Reference Model [10]. This is an application profile "to provide a comprehensive suite of e-learning capabilities that enable interoperability, accessibility and reusability of Web-based learning content."

IEEE LOM: IEEE 1484 Learning Objects Metadata. A Learning Object is defined as any entity, digital or non-digital, which can be used, re-used or referenced during technology-supported learning [11]. The mapping of LOM to Dublin Core is available at [12].

IMS Global Learning Consortium. IMS Meta-Data v1.2.2 [13] The IMS initiative originated in higher education but it now involves stakeholders in corporate and government training, K-12, and continuing education.

The above schemas are interwoven. For example, SCORM uses LOM vocabularies. All explicitly map to qualified Dublin Core elements. Educational objects in general are concerned with usage or the re-purposing of content, not necessarily with storage. Extensions to our model (for example, the addition of the Audience element) as well as value lists (element types) allow for incorporating some of these needs.

3. Why develop PB Core?

Many parties have asked us why we did not adopt and adapt metadata schemas already in existence or in development. For several reasons, the existing standards were not appropriate to our needs. Basically, alternative schemas were either too cursory in their descriptive capabilities or far too ponderous.

An implementation project, such as the Public Broadcasting Metadata Initiative Project, generally finds that no one metadata standard completely meets its needs for descriptions of media essence. General standards, like Dublin Core, are often folded into domain- or sector-specific standards, such as MPEG-7 for multimedia and IEEE/LOM for educational resources. New elements may be devised which meet local needs not covered by any existing standards. The Public Broadcasting Core can be thought of as an application profile whose schema combines elements from multiple standards, with application-specific constraints (as in the use of specific controlled vocabularies or structured values). The PB Core must be understandable and usable by all public broadcasting entities, from the smallest local NPR radio station to the largest public television producers of national programming.

The PBMI's primary interest is in data exchange, data crosswalks, and interoperability, not necessarily in creating a complete metadata model that can be exploited by digital asset management systems for comprehensive, original cataloging and markup of essence. The Project desires to facilitate the sharing of metadata and the discovery of valued assets. The PB Core is intended to be "simple," but not "simplistic." Furthemore, the PB Core should be considered as a starting point that may accommodate metadata extensions of interest to specific communities and users.

Consequently, the Project undertook a path that would reflect the Public Broadcasting industry's needs and wants regarding media assets by gathering together representatives from public broadcasting and growing a consensus. The unique quality of public broadcasting, both television and radio, is its local ownership and local ties to its surrounding communities. In a parallel fashion, the Public Broadcasting Metadata Initiative was designed to tap into the various local constituencies and develop a metadata core from "grassroots" origins, rather than by administrative edict.

The Project conducted a detailed "needs assessment" of public broadcasters. Such measures are revealing and often unmask and articulate conditions, issues, needs, and desires that otherwise are dismissed or forgotten. By applying user-centered techniques PBMI was able to discover a wide spectrum of needs and applied the most appropriate metadata elements.

4. The Process for Assessing the Need and Gathering User Requirements

Public broadcasters have always endeavored to engage in complex and robust relationships with their constituents, whether those are viewers, listeners, educators, community leaders, etc. We have always provided extensive outreach for our broadcast content, with particular emphasis on the needs of K-12 teachers and lifelong learners. Today, with the advent of the Internet, that outreach is more significant and successful than ever before. As mentioned above, we also have an extremely complex structure; as opposed to our media counterparts, who increasingly concentrate their ownership and control of media outlets, very little of public broadcasting's operations are centralized. We have innumerable systems for producing and tracking our content, and our institutions are structured in a variety of ways, often based on who holds the broadcast license.

In order to ascertain the metadata needs of our "external" users – constituents – and "internal" users – local and national staff – we first created a straightforward list of users (see appendix), and then double-checked this "strawman" with the core PBMI working group. A "User Requirements Team" was formed from within the working group. Using the now-modified user list, they set out to create a series of Use Case Scenarios. The team wanted to find out during which present or future metadata data search, delivery, or translation activities users would encounter and rely on standardized protocols for content descriptions.

During this process, the "User Requirements Team" interviewed a large number of stakeholders, including national program distributors, local station broadcast operations and IT staff, a K-12 "learning object" consortium, an independent television production company, a television graphic artist, and "interactive" specialists (web and TV).

What emerged was a clear division between full-program metadata (such as title, format, date), which serves the needs of national distribution and local broadcast operations, and fragment, or clip-level data, which serves the needs of producers, educators, and website programmers.Most use case participants felt that it was critical to have a simple, intuitive set of metadata elements, with extensions for particular constituencies, e.g., K-12 curriculum-correlation, or graphics creation,, so that the maximum number of assets could be identified and retrieved by the greatest number of individuals and institutions.

There was a great deal of concern about rights management, without which future business and service models crumble. Several interviewees felt that the working group should also determine standards for metadata exchange, such as XML.

5. The Process of Refining the PB Core

A powerhouse of motivated and opinionated experts were assembled to contribute to the Public Broadcasting Metadata Initiative Project. The members were drawn from a variety of communities related to public broadcasting:

- National public television organizations and program distributors
- National public radio organizations and program distributors
- Online Internet-based resource organizations
- National program producers
- State and regional network organizations
- Community radio and TV licensed stations

- University radio and TV licensed stations
- Educators
- Metadata subject matter experts

The initial work of the members for the Public Broadcasting Metadata Initiative Project lasted seven months. The overarching goal of the group was to recommend usable metadata fields that would facilitate the exchange of program and resource information between public broadcasting communities and other interested parties. Guiding our work process was the question, "How would a particular metadata element ultimately contribute to the discovery of public broadcasting's intellectual content by various end-users"? The objectives of the Working Group were to:

- Develop and refine user requirements for a sharable metadata element set.
- Review existing metadata schemas to determine their applicability to the public broadcasting arena, to identify gaps and overlaps, and to incorporate the most germane while discarding the least useful or confounding.
- Determine the scope and breadth of a usable metadata schema that was consensus-built, extensible, and interoperable with other asset management systems and databases.
- Draft a preliminary application profile of the public broadcasting core metadata of descriptors and their usage.
- Present the PB Core to the public broadcasting community for review and comment.
- Refine and revise the PB Core prior to release and publication.

In the seven-month time period, two full meetings of the entire Working Group were conducted, as well as follow-up committee work.

- First Meeting: 2002-4-24&25
- Committee Work: 2002-Summer, PB Core Review Team and User Requirements Team
- Second Meeting: 2002-9-12&13
- The Boston Summit: 2002-10-16,17,18

These activities led to an intensive three-day work session in Boston (2002-10-16,17,18), where the Public Broadcasting Metadata Core was refined and honed by the PB Core Review Team.

Before the Boston Summit, the PB Core Review Team had surveyed existing metadata dictionaries from various authorities and organizations, including those in use by several public broadcasting groups. A total of 467 separate metadata elements were compiled, which spawned 2335 recommendations for grouping and collapsing the elements into the most relevant. From these recommendations, a total of 249 working metadata elements and their qualifiers were selected.

The work of the PB Core Review Team at the Boston Summit combined redundant elements, discarded the less relevant, and debated the appropriate application of preferred metadata within the dictionary. The Summit yielded a preliminary draft of 59 metadata elements and their qualifiers that were most appropriate to public broadcasting and related communities.

6. The Public Broadcasting Core Elements

Many of the 59 metadata elements selected for the Public Broadcasting Core of metadata descriptors were drawn from the Dublin Core Metadata Initiative. Others were retained from existing public broadcasting digital asset management systems in development. Still others were drawn from additional working groups.

Table 1 reviews the 59 elements and qualifiers currently under consideration by the Public Broadcasting Metadata Dictionary Project. The Registration Authorities listed represent the agency of responsibility for the long term integrity and viability of particular metadata elements and associated qualifiers:

- DCMI: Dublin Core Metadata Initiative
- DC-Ed: DCMI Education Working Group
- ViDe: Video Development Initiative
- [PBCore]: Corporation for Public Broadcasting as Interim Steward
- [MPR]: Minnesota Public Radio as Interim Steward

Table 1. Recommended Metadata Elements of the Public Broadcasting Metadata Initiative Project

Element Name	**Registration Authority and Element Definition**
01.00 Title	DCMI: A name given to a resource, as well as any other title(s) that would be useful in uniquely identifying a resource and that would facilitate discovery and retrieval.

Element Name	Registration Authority and Element Definition
01.01 Title.Alternative	DCMI: An Alternative Title is used in order to identify an asset or resource that has a title similar to the proper title, but which further assists in discovery and retrieval.
01.02 Title.Series	[PBCore]: A Series Title is one specifically identified by the video or audio production agency and is named as such in order to facilitate discovery and retrieval, as well as to more accurately reflect how a resource's title fits into a hierarchy of proper titles that are used to describe it.
01.03 Title.Episode	[PBCore]: An Episode Title is one specifically identified by the video or audio production agency and is named as such in order to facilitate discovery and retrieval, as well as to more accurately reflect how a resourceis title fits into a hierarchy of proper titles that are used to describe it.
01.04 Title.Program	[PBCore]: A Program Title is one specifically identified by the video or audio production agency and is named as such in order to facilitate discovery and retrieval, as well as to more accurately reflect how a resource's title fits into a hierarchy of proper titles that are used to describe it.
02.00 Creator	DCMI: An entity primarily responsible for making the content of the resource or asset. May be a person, business, organization, group, initiative or service.
02.01 Creator.Role	[PBCore]: Unlike print resources, there is no single role, such as author, that is commonly understood to have primary responsibility for the intellectual content of many resources, such as audio, video or film assets. In such cases, creators can include many different roles deemed to have primary responsibility for the creation of the essence, such as the instructor for a video course, the interviewee from a video history program, or the director of a feature film.
03.00 Subject	DCMI: The topic(s) of the intellectual content of a resource or asset. Contains controlled values and uncontrolled values (keywords). Use the Description element for more free-form text descriptions of a resource.
04.00 Description	DCMI: An account of the intellectual content of the resource. Descriptions are more free-from text entries when compared to the controlled vocabularies associated with the Subject element.
04.01 Description.Abstract	DCMI: As an account of the content of the resource, the qualifier Abstract is a short narrative summary of the topic of the resource. Provides additional supplied text by experts that adds color or insight to the description of the resource or asset not otherwise identified in the more specific content related fields. Anecdotal comments welcomed.
04.02 Description.Table of Contents	DCMI: As an account of the content of the resource, the qualifier Table of Contents is used for partial or full listings of subunits of the resource. Use the Table of Contents to identify other descriptive information such as: Composers and Works contained in a program; Cue Sheets; Play Lists; Rundowns; Edit Decision Lists (EDLs) (unformatted); Content Flags; Index of Sections or Segments; Formal Table of Contents.
04.03 Description.ProgramRelatedText	[PBCore]: As an account of the content of the resource, the qualifier ProgramRelatedText identifies other audio and textual representations of the main audio or language presentation mode for a resource or asset.
05.00 Publisher	DCMI: An entity responsible for distributing or making a resource available to other end-users and communities. May be a person, business, organization, group, initiative or service. Some resources may not have a publisher or distributor, and thus will not have an entry under Publisher.
05.01 Publisher.Role	[PBCore]: The Role that is played by a specific Publisher or Publishing entity is identified.
06.00 Contributor	DCMI: An entity responsible for making contributions to the content of the Resource, but whose contribution is secondary to any entity specified in the Creator element (for example, film editor, screenwriter, narrator). Examples of Contributor include a person, an organization, or a service. Typically, the name of a Contributor should be used to indicate the entity.

Element Name	Registration Authority and Element Definition
06.01 Contributor.Role	[PBCore]: The Role which a Contributor plays is identified here. Use this element and qualifier to identify important production credits for a resource, e.g., producer, director, writer, special thanks, funding agencies, programmers, designers, graphics, instructional design, etc.
07.01 Date.Created	DCMI: The creation date for a resource or program.
07.02 Date.Issued	DCMI: Date of formal issuance (e.g. publication) of a resource for general public consumption.
07.03 Date.AvailableStart	DCMI: A specific start date for a resource's availability. May refer to start dates for the availability of a program that is broadcast locally, regionally, nationally or internationally.
07.04 Date.AvailableEnd	DCMI: A specific date that a resource's availability has come or will come to an end. May refer to end dates for the availability of a program that is broadcast locally, regionally, nationally or internationally.
08.00 Type	DCMI: The nature or genre of the content of the resource, or the purpose for which the asset was created and made available.
08.01 Type.Form	[PBCore]: A format or program category for a resource.
08.02 Type.Genre	[PBCore]: The nature or genre of the content of a resource.
09.01 Format.Physical	[PBCore]: A physical manifestation of a resource as it may exist as a format or carrier that occupies physical space dimensions.
09.02 Format.Digital	DCMI: A digital instantiation of a resource that may or may not have existed originally in an analog, physical form. Digital media formats may be expressed as formal Internet MIME types or as other means of expressing the format of a digital resource.
09.03 Format.Identifier	[MPR]: Identifying information about the format of a resource.
09.04 Format.FileSize	ViDe: Measures the storage requirements or file size of a digital resource in Bytes, Kilobytes, Megabytes or Gigabytes to provide the most meaning to the end user.
09.05 Format.AudioBitDepth	[PBCore]: For a program or resource, this qualified element measures an audio signal in a number of bits and answers the question, 'How Much' data is allocated to a digital sampling of an audio signal. Provides information important for identifying retrieval and playback/display requirements for a resource.
09.06 Format.AudioChannelConfiguration	[PBCore]: Indicates the number of audio channels configured for the playback of a resource.
09.07 Format.AudioDataRate	[PBCore]: Expressed as amount of data per second and indicates how much data is delivered through a particular delivery pipeline for every second.
09.08 Format.AudioSamplingRate	[PBCore]: Measured in kiloHertz for a program or resource, this qualified element quantifies 'How Much' data is allocated to a digital sampling of an audio signal. Provides information important for identifying retrieval and playback/display requirements for a resource.
09.09 Format.ImageAspectRatio	[PBCore]: Indicates the ratio of horizontal to vertical proportions in the display of an image or moving image.
09.10 Format.ImageBitDepth	[PBCore]: For a program or resource, this qualified element measures a still or moving image in terms of the number of bits in a sample, and answers the question, How Much data is allocated to a digital sampling. Provides information important for identifying retrieval and playback/display requirements for a resource.
09.11 Format.ImageChannelConfiguration	[PBCore]: Indicates the number of image channels available in a resource. May be most appropriate for digital files, like QuickTime in which multiple video tracks can be encoded in a single file.

Element Name	Registration Authority and Element Definition
09.12 Format.ImageColorCode	[PBCore]: Indicates the color or lack of color in an asset. Does not measure the specific color metrics of a image or moving image.
09.13 Format.ImageDataRate	[PBCore]: Expressed as amount of data per second and indicates how much data is delivered for an image or moving image through a particular delivery pipeline for every second.
09.14 Format.ImageFrameRate	[PBCore]: Indicates the frames per second found in a resource's playback or display.
09.15 Format.ImageFrameSize	[PBCore]: Indicates the horizontal and vertical resolution of a format type. May be expressed in pixels, pixels per inch, or in the case of ATSC digital TV, a combination of pixels measured horizontally vs. the number of lines of image data stacked vertically (interlaced and progressive scan).
09.16 Format.TimeStart	[PBCore]: Indicates a time stamp representing the beginning point for the playback of a resource. Use in combination with Format.Duration to identify a sequence or segment of a resource that has a fixed start time and end time.
09.17 Format.Duration	ViDe: Describes the duration in time units for a resource, if that resource has an identifiable, linear start-to-end playback. Format.Duration does not describe the time required to utilize a resource in a setting, but is rather a strict playback time, TimeStart to TimeEnd.
09.18 Format.Standard	[PBCore]: The standard refers to an overarching architecture for underlying media formats.
09.19 Format.Type	[PBCore]: The Qualifier of Type is hierarchically a subset of the values found under Format.Standard and describes specific kinds of media formats found for each media standard.
09.20 Format.Encoding	[PBCore]: This proposed element with qualifier is designed to offer a single element with which the various media standards and their collected format types can be identified for a particular resource.
10.00 Identifier	DCMI: An unambiguous reference or identifier for a resource within a given context. Best practice is to identify a resource by means of a string or number corresponding to an established or formal identification system.
11.00 Source	DCMI: A reference to another resource from which the present resource is derived.
12.00 Language	DCMI: The primary language of the intellectual content of the resource, usually expressed by the audio track. If other, alternative audio and textual representations of the main audio or language presentation mode exist for a resource or asset, describe that information in the Language.Usage element.
12.01 Language.Usage	[PBCore]: The qualifier Language.Usage identifies the existence of other audio and textual representations of the main audio or language presentation mode for a resource or asset.
13.01 Relation.Type	[PBCore]: Relation.Type identifies a second resource that is related to the primary resource. It defines the relationship between the second resource and the primary resource. While the primary resource is described by the rest of the asset management's database record, the second resource is described using the Relation field.
13.02 Relation.Identifier	[PBCore]: Identifies a second resource related to the primary resource by using a specific numbering or labeling scheme to call out the related resource. Used in combination with the Relation.Type element to cross reference the type of relation with a unique identifier for that relation.
14.01 Coverage.Spatial	DCMI: Identifies the extent or scope of the resource's content from a spatial or geographical perspective of the intellectual content of a resource. Coverage.Spatial is used for geographic coordinates of maps and map-like images (e.g. aerial maps or map-like images concatenated as a video file) or to associate place names or logical jurisdication for a resource.

Element Name	Registration Authority and Element Definition
14.02 Coverage.Temporal	DCMI: Identifies the extent or scope of the resource's content from the perspective of the temporal or time characteristics of the intellectual content of a resource.. CoverageTemporal is used for date and time-based events, designated numerically for precision searching, where the time element is critical for identification and use of the resource.
15.00 Rights.Usage	[PBCore]: Information about rights held in and over the resource, particularly in what manner the resource will be used, eg., broadcast, web, PDAs, or education/classroom.
15.01 Rights.Reproduction	ViDe: Statements or references about rights held in and over a resource, specifically regarding the rights to reuse, repurpose or reproduce a resource.
15.02 Rights.Access	ViDe: Access information about rights held in and over a resource. Rights.Access indicates either 'open access' or 'restricted access.' These two options are used as flags to trigger certain actions. For example, metadata records with 'restricted access' will not be exposed for mining by OAI initiatives.
16.01 Audience.Level	DC-Ed: A general statement describing the education or training sector. Alternatively, a more specific statement of the location of the audience in terms of its progression through an education or training sector or level.
16.02 Audience.Rating	[PBCore]: Designates categories of users for whom the resource is intended or judged appropriate. Standard ratings have been crafted by the broadcast television industry which are useful.
17.00 Standard	DC-Ed: A reference to an education or training standard with which the resource is associated. These standards may be designated at a national level, state level or local level. Often the standards are represented by topical areas, course of study, learning objectives, or locally significant identification schemas and numbering systems.
18.00 Annotation	[PBCore]: General field to be used to append helpful information for the metadata markup team about an asset and its metadata.
19.00 Location	[PBCore]: Identifies the location of a specific format or instantiation of a resource. Usually a text string describing where in the physical world the resourceís physical format resides. This is not an ID number, but a location identifier. Used when the actual physical form is being indexed in an asset management system.

7. The Application Profile

The 59 elements are delineated by 15 attributes according to the modified ISO 11179 Specification and Standardization of Data Elements [14]. .The full accounting of the specification is too large a document to include in this paper.

1. Element Number
2. Element Name
3. Version of the Element
4. Element Label
5. Definition
6. Namespace Identifier
7. Registration Authority
8. Language of the Element
9. Obligation in Usage
10. Data Type
11. Maximum Occurrence
12. Encoding Schemes
13. Restricted Values
14. Examples
15. Usage Guidelines

PBMI's interest is in data exchange, data crosswalks, and interoperability, not necessarily in creating a complete metadata model that can be exploited by digital asset management systems for comprehensive, original cataloging and markup of essence. Consequently, the primary desire of PBMI is to facilitate the sharing of metadata and the discovery of valued assets. Within the Application Profile, issues of concern to PBMI are:

- Who will serve as the real registration authority that takes responsibility for the declaration and maintenance of our newly defined, custom elements and their qualifiers that are not already part of a standard?
- Who will publish versions of the Core and its updates? Provide documentation? Provide Technical Support?
- How will we monitor adoption and compliance?
- How will we measure successful implementation?

The Project recognizes that it needs to remain focused on the fact that the Working Group is not a body of "standards makers." Rather, we are "real life implementers" who are tasked with generating effective solutions in order to service the efficient and widespread delivery of public broadcasting's intellectual content. Similar to our day-to-day business, we are engaged in applied and practical solution-making.

Like many other groups debating the application of metadata schemes, the Project remains conflicted in how best to match metadata descriptors with various instantiations of essence and assets. The question of embracing a "one-to-one" relationship between a metadata record and its associated essence or subscribing to a "one-to-many" relationship between a metadata record and the various instantiations of its essence still plagues the PBMI Project. Compelling arguments have been presented on both sides of the issue. We are hopeful that the next phase of our project, a Request for Comments, will assist us in sorting out a solution.

8. Feedback and Evaluation Mechanisms

To a great extent, the work of the Public Broadcasting Metadata Working Group has modeled an unheard-of process – coordination and consensus across vastly different institutions, on a topic of extreme detail and importance. The Preliminary PB Core is ready to be reviewed and tested.

During the next several months the Working Group will be asked to engage in an even more difficult process – a mid-course evaluation.

The group will be divided into task teams, and through research, interviews, conference calls, and "thought papers," will address the following issues and objectives:

- determine that the PB Core is sustainable over time (including knowing its lifespan, form, cost, etc. and how stations and producers can be made to comply with the protocol);
- ensure that the PB Core's strategic value is understood and acknowledged by senior management who will need to support it;
- devise a plan to integrate the PB Core into the day-to-day operations within local and national content infrastructures, especially the PBS Next Generation Interconnection System and NPR's Content Depot.

It is our assumption that these difficult questions will be answered in a manner that leads the project to the RFC (Request for Comments) process, and then test implementations in typical metadata scenarios.

The RFC process will include other public broadcasting production, IT and broadcast operations staff, key software vendors serving the industry, standards organizations, partnering institutions, etc.

Test implementations of the PB Core, still to be determined, will likely include radio, television and website production collaborations, tape libraries, national program distribution, etc.

9. Next Steps

The PBMI process has illuminated for participants and observers alike the critical need for a new, "advanced networking" approach toward conducting our core activities. We must change our institutions and infrastructures, even our funding models, to reflect a new spirit of exchange, collaboration and consolidation. Certainly, without Internet-like standards for descriptive and administrative metadata, rich media file formats, file exchange, etc., we will not be able to keep pace with changes in the media environment, nor will we advance our public service mission.

References

[1] Dublin Core Metadata Element Set, Version 1.1: Reference Description. Retrieved April 15, 2003, from http://www.dublincore.org/documents/dces/

[2] OAIS: Reference Model for an Open Archival Information System. Retrieved April 15, 2003, http://ssdoo.gsfc.nasa.gov/nost/isoas/overview.html

[3] SMEF-DM: Standard Media Exchange Framework - Data Model. Retrieved April 15, 2003, http://www.bbc.co.uk/guidelines/smef/

[4] MPEG-7: "Multimedia Content Description Interface" ISO/IEC JTC1/SC29/WG11. Retrieved April 15, 2003, from http://mpeg.telecomitalialab.com/standards/mpeg-7/mpeg-7.htm

[5] Hunter, J. "A Proposal for the Integration of Dublin Core and MPEG-7", ISO/IEC JTC1/SC29/WG11 M6500, 54th MPEG Meeting, La Baule, October 2000. Retrieved April 15, 2003, from http://archive.dstc.edu.au/RDU/staff/jane-hunter/m6500.zip

[6] Hunter, J. ViDE Video Access Group, "An Application Profile which combines DC and MPEG-7 for Simple Video Description", February 12, 2002. Retrieved April 15, 2003, from http://archive.dstc.edu.au/RDU/staff/jane-hunter/publications.html

[7] Agnew, G. (2003) A Tale of Two Schemas: Mapping Dublin Core to MPEG7. Retrieved May 15, 2003, from http://gondolin.rutgers.edu/MIC/text/how/mpeg7DC_agnew.pdf

[8] MARC: Machine Readable Cataloging. Retrieved April 15, 2003, from http://www.loc.gov/marc/

[9] METS: Metadata Encoding and Transmission Standard. Retrieved April 15, 2003, from http://www.loc.gov/standards/mets/

[10] SCORM: The Sharable Content Object Reference Model. Retrieved April 15, 2003, from http://www.adlnet.org/

[11] IEEE Learning Technology Standards Committee (LTSC). Learning Object Metadata. Draft Document v3.6. Retrieved April 15, 2003, from http://ltsc.ieee.org/doc/wg12/LOM3.6.html

[12] Sutton, S.A. (1999) IEEE 1484 LOM mappings to Dublin Core: *Learning Object Metadata: Draft Document v3.6*, IEEE Learning Technology Standards Committee (LTSC), 5 September 1999. Retrieved April 15, 2003, from http://www.ischool.washington.edu/sasutton/IEEE1484.html

[13] IMS Global Learning Consortium. IMS Meta-Data v1.2.2. Retrieved April 15, 2003, from http://www.imsproject.org/metadata/index.cfm

[14] ISO 11179 Specification and Standardization of Data Elements. Retrieved January 11, 2003, from http://www.diffuse.org/meta.html#ISO11179.

Appendix: Public Broadcasting Metadata Initiative Project Participants

At Large

Marcia Brooks, WGBH (Project Director)
Dennis Haarsager, KWSU
Amy Rantanen, WGBH
Alison M. White, Corporation for Public Broadcasting

Advisors

Grace Agnew, Rutgers University (AMIA)
Judy Brown, University of Wisconsin, DOD Academic CoLab (SCORM Project)
Efthimis N. Efthimiadis, The Information School, University of Washington, Seattle, WA

Working Group

Alan Baker, Minnesota Public Radio
Nancy Baldacci, American Public Television
Sharon Blair, AMIA Local Television Task Force
Marty Bloss, National Public Radio Public Radio Satellite Service (NPR PRSS)
Paul E. Burrows, Media Solutions, OIT, University of Utah
Brian Callahan, WHRO
Michael Connet, onCourse
David Felland, Milwaukee Public Television
Tom Handy, KWSU
Steven Heard, Public Interactive
Rob Holt, NPR Online
Dave Johnston, PBS Online
Ann Lootens, WGBH
Dave MacCarn, WGBH
Chuck McConnell, NETA/OSBE
Bea Morse, Public Broadcasting Service (PBS)
Robin Mudge, onCourse (*former participant)*
Lesley Norman, David Grubin Productions
Meg O'Hara, WNET
Tim Olson, KQED
Marilyn Pierce, Public Broadcasting Service (PBS)
Richard Ruotolo, Public Radio International (PRI)
James Steinbach, Wisconsin Public Television
Brent Trinacty, Public Radio International (PRI)*
Cate Twohill, Public Broadcasting Service (PBS)
Steven Vedro, Consultant
Tracy Vosburgh, WPSX (Penn State University)
Art Zygielbaum, Nebraska Education Television

Additional Support Provided By

Scott Bridgewater, National Public Radio Public Radio Satellite Service (NPR PRSS)
Carrie Lowe, Public Broadcasting Service (PBS)
Thom Shepard, WGBH
and other staff at many of the participating organizations listed above

ATSC Datacasting: Opportunities and Challenges

Gomer Thomas
LG Electronics Research Center of America, Inc. (LGERCA)
Princeton Junction, NJ

Abstract

ATSC datacasting provides many opportunities for terrestrial DTV broadcasters. This paper discusses key issues in this exciting area. Topics include:

- A taxonomy of data broadcasting in terms of target audience and application characteristics
- An overview of the emerging ATSC Data Broadcast Standard
- Challenges in implementing end-to-end data broadcast solutions for enterprise-to-enterprise applications

One key challenge is managing the end-to-end flow of data, with suitable architectural support for content providers, broadcasters, and users. Other challenges include bandwidth allocations, error correction, compression, and security.

Introduction

Terrestrial TV broadcasters are being forced to bear heavy expenses in making the conversion to digital television (DTV) broadcasting. It is not at all clear that they will be able to recover these costs through their normal revenue sources (primarily advertising). They may be able to offer more channels of TV programming than before, since DTV enables the packing of multiple standard definition (SDTV) channels into a single 6 MHz broadcast band, but they will still be competing for more or less the same number of viewer eyeballs and advertising dollars.

Data broadcasting (or "datacasting" for short) offers the potential for entirely new revenue sources. With modern encoders, even a high definition (HDTV) broadcast only requires 16 to 18 Mbps of digital bandwidth, out of the entire 19.38 Mbps available in a 6 MHz broadcast band under the ATSC standard. This leaves a significant amount of bandwidth that can be used for arbitrary data.

There are basically two ways in which broadcasters can use this leftover bandwidth to generate additional revenue:

- Use the excess bandwidth to broadcast data which enhances the appeal of their TV programming and/or TV advertising in an attempt to attract more advertising dollars.
- Lease the excess bandwidth to other enterprises who want to distribute data to large numbers of users in the broadcaster's viewing area.

In practice, both of these approaches will likely be used to varying degrees by different broadcasters.

This paper first analyzes the opportunities that arise from both of these approaches, and describes the mechanisms defined in the ATSC Data Broadcast Standard which can be used to support these opportunities. It then focuses on the option of leasing bandwidth to other enterprises, analyzing the challenges which arise in implementing end-to-end solutions, with special attention to system architecture.

Application Taxonomy

There are a number of axes which can be used to classify datacasting applications. The single most important of these in terms of both the business model and the technological infrastructure is whether the broadcast data is targeted to enterprises or to the consumer mass market.

Printed from a paper delivered at the NAB Broadcasting Engineering Conference, Apr. 8-13, 2000.

Consumer Market. For datacasting applications targeted to consumers, a key requirement for success is that large numbers of consumers have a DTV receiver (DTV set, DTV set-top box, or PC with DTV card) which can receive and use the broadcast data. For this to happen:

1. There must be substantial penetration of DTV receivers into the consumer market.
2. The datacasting standards must be reasonably mature, so that they can be supported by ordinary, off-the-shelf DTV receivers.

It may take a while yet for these conditions to be met. In the meantime, consumers are unlikely to pay substantial sums for specialized receivers to access broadcast data.

Utilizing broadcast data requires not only standards for encoding the data in the broadcast stream, but also standards for applications in the receiver to operate on the data. This could be in the form of specifications for one or more standard applications in the receiver, such as a standard HTML engine. Alternatively, it could be in the form of specifications for a standard execution environment in the receiver, so that a variety of applications could be downloaded from the broadcast stream and executed.

The ATSC technical subgroup on the DTV Application Software Environment (DASE) is working on a standard that incorporates both these elements. However, it may be some time before it is fully approved and supported by consumer products.

The standardization requirement can be relaxed somewhat for applications targeted to consumers with PCs and DTV cards, since custom software to support such applications can be downloaded from a Web site. However, such applications are still only feasible when large numbers of consumers have DTV cards, which may take a while.

Datacast applications targeted to consumers can be further classified by the degree of coupling to the normal TV programming, as shown in Table 1.

Tightly coupled data are intended to enhance the TV programming in real time. The viewer tunes to the TV program and gets the data enhancement along with it. In many cases the timing of the display of the broadcast data is closely synchronized with the video frames being shown.

For example, a sports event enhancement may allow the viewer to call up and dismiss the time clock on command. When the clock is showing, the time on the clock must be closely synchronized to the action on the screen.

Loosely coupled data are related to the program, but are not closely synchronized with it in time. For example, an educational program might send along in the broadcast some supplementary reading materials or self-test quizzes. These might not even be viewed at the same time as the TV program. They may be saved in the DTV receiver and perused later.

Non-coupled data are typically contained in separate "data-only" virtual channels. They may be data intended for real-time viewing, such as a 24-hour news headline or stock ticker service, or they may be data intended for use completely outside the DTV context.

An example of the latter usage would be delivery of very large files, such as MP3 music files or PC software files. A consumer might purchase the right to receive such files from a Web site. The files would be inserted into a DTV broadcast in encrypted form. The purchasers would receive keys which could be used to decrypt them.

Coupling to Normal TV Programming	Sample Applications
Tight (program enhancement)	Statistics during sports event Story details during newscast Interactive component to ads
Loose (program augmentation)	News and/or stock ticker during newscast Supplementary materials for educational broadcasts
None	24-hour stock ticker service Delivery of on-line music sales Delivery of software updates

Table 1. Datacast Applications for Consumers

For most consumer applications the broadcaster's revenue would come from advertising. However, in the case of non-coupled data intended for use outside the DTV context, the broadcaster might be paid for the bandwidth by a third party that actually drives the business. In an e-commerce situation, the payment may be calculated as a percentage of sales.

Enterprise-to-Enterprise Market. For datacast applications targeted to enterprises, there is not so much need for DTV market penetration and standardization. The value of such applications to enterprise customers may easily be large enough to justify the cost of specialized receivers. Thus, enterprise-to-enterprise applications may be viable from an economic standpoint much sooner than consumer applications.

Almost all enterprise-to-enterprise applications will involve data which is not coupled to the regular TV programming.

Examples of enterprise-to-enterprise applications could include:

- Distributing media-rich real estate listings to local real estate offices
- Distributing stock quotes, company reports, industry analyses, etc., to day trading offices
- Distributing tourist information and related advertising to kiosks in hotel lobbies, shopping malls, tourist centers, etc.
- Distributing educational materials, including pages from selected web sites, to public schools
- Distributing information among state and local government agencies, such as agricultural bulletins and police reports
- Distributing audio clips, such as traffic, weather, and news reports, from a TV station to the local radio stations for use in their broadcasts
- Broadcasting training courses to commercial customers, in the form of low-bandwidth video and ancillary materials

Applications such as these may be practical now.

Data Broadcast Standard

The emerging ATSC Data Broadcast Standard [1] contains specifications for three separate aspects of data broadcasting:

- Data broadcast protocols/formats
- Data announcement protocols
- Data discovery and binding protocols

The data broadcast protocols specify the precise formats for inserting data into the broadcast stream.

The data announcement protocols specify how to put information in the broadcast stream which can be used by receivers to inform viewers of upcoming broadcasts of data programs or data-enhanced programs.

The data discovery and binding protocols tell how to put information in the broadcast stream which will allow applications running in receivers to identify and locate broadcast data in a program.

Data Protocols. The ATSC Data Broadcast Standard [1] supports the broadcast of 9 basic types of data:

- Synchronized streaming data
- Synchronous streaming data
- Asynchronous streaming data
- Synchronized data modules
- Asynchronous data modules
- Synchronized multiprotocol datagrams
- Synchronous multiprotocol datagrams
- Asynchronous multiprotocol datagrams
- Data piping

Synchronized data has a strong timing association with another stream, such as a video stream. Each data item in a synchronized data stream is intended to be presented to the viewer at a specific time relative to a clock established in the other stream, for example at the same time as a specific video frame in a video stream.

Synchronous data has an internal timing association among its data items, but no timing association with data in another stream. The presentation time of the data items relative to each other is specified, but not the time relative to anything else in the broadcast stream.

Asynchronous data has no internal or external timing requirements.

Streaming data will typically be used in applications which require streaming audio and/or video, but which for various reasons do not want to encode the audio or video in one of the ATSC standard data formats. An example might be low bandwidth audio or video, or perhaps MPEG-4 content.

A typical use of data modules will be to transmit files. Asynchronous data modules are typically broadcast cyclically, in the form of so-called "data carousels," so that the receiver will have access to them regardless of the point in the program at which the viewer tunes to the channel.

The Data Broadcast Standard supports multiprotocol datagrams for pretty much any protocol, but it is widely expected that the most common usage will be for IP datagrams. In fact, the standard contains some facilities which simplify the handling of that special case.

Data piping is intended to support proprietary datacasting applications. It simply specifies that data will be packed into the payload of the MPEG-2 transport packets and leaves it completely up to the application how the data is formatted inside the packets.

The term "data service" is used to refer to any collection of one or more data types.

Data Event Announcements. The Data Broadcast Standard [1] defines "Data Event Tables" (DETs) for data services, analogous to the "Event Information Tables" (EITs) for audio/video programs. An entry in a DET may describe a data-only program, or it may describe the data component of an audio/video/data program in cases where the broadcaster wishes to announce the data service separately.

For example, a broadcaster may include data to support an interactive stock market quote service in with a financial news program. In that situation the broadcaster may have an entry in the EIT for the news program and a separate entry in the DET for the associated stock market query service. Both entries would reference the same virtual channel and same time slot. The virtual channel would contain video streams, audio streams, and data streams, which the DTV receiver would sort out with the help of a PSIP table called the "Virtual Channel Table."

Another broadcaster may have a more sophisticated interactive stock query service which provides general economic forecasts by industry, background analysis reports on individual companies, and stock market quotes. This may run as a stand-alone service in its own data-only virtual channel. Such a service would be announced only by an entry in the DET, since it does not have an audio/video component.

Yet another broadcaster may enhance the broadcast of a sports event with an interactive statistics service, so that a viewer can see game statistics, biographical material on the competitors in the event, etc., during the event. Such an event would typically have an entry only in the EIT, since the data service is really an inseparable part of the sports event.

The standard also defines a "data broadcast descriptor" which can be associated with events in EITs or DETs. As well as announcing the presence of a data service, the data broadcast descriptor gives information about the service which can be used by the receiver to determine how much buffer space to allocate for the service, etc.

It is expected that data-capable DTV receivers will have interactive electronic program guides (EPGs) which show the data services defined in the DETs as well as the ordinary TV programs defined in the EITs. Moreover, TV programs which are data-enhanced, such as the sports event described immediately above, will typically be highlighted in the EPG so that viewers can identify data-enhanced programs. Many viewers may choose to watch them in preference to non-enhanced programs, since they are likely to provide a richer viewing experience.

Data Discovery and Binding. The Data Broadcast Standard [1] defines a "Data Service Table" (DST) that is to be included in the broadcast of any program which contains a data service. The DST lists all the data included in the data service, telling where to find the data items, what data type they are, and what to do with them.

For example, a data service might consist of a collection of HTML pages carried in a data carousel. The DST would tell where in the broadcast stream to find the data carousel, would indicate that the modules are HTML pages, and would tell which page is the top-level page that should be displayed to start up the data service.

In some cases a data service might have components that are not being broadcast in the virtual channel itself. Some components might be in other virtual channels. In the case of broadcasts intended for DTV receivers with Internet connections, some components might even be files on Internet servers. The Data Broadcast Standard defines a "Network Resource Table" (NRT) that can be used to identify and locate such remote components. The DST points to entries in the NRT, and the NRT provides complete information on the location of the remote data.

Implementation Challenges

The remainder of this paper focuses on enterprise-to-enterprise datacasting applications (sometimes called Business-to-Business or B2B applications), since these appear to offer better near term revenue opportunities.

There are a number of challenges in implementing end-to-end solutions for enterprise-to-enterprise data broadcast applications. Key considerations include:

- Business roles of the key players
- Data flow management
- Bandwidth management
- Security
- Error prevention and recovery
- Compression

Roles. There are three primary players involved in enterprise-to-enterprise datacasting:

- Data Provider (or Service Provider)
- Broadcaster
- Data Subscriber

A Data Provider is an enterprise which provides data to be broadcast. It may generate the data itself, or it may aggregate the data from other sources. The Data Provider is typically the direct customer of the Broadcaster, paying the Broadcaster for the bandwidth used to transmit the data.

A Data Subscriber is an enterprise which receives data and uses it. The Data Subscriber is typically the direct customer of the Data Provider, paying for the data it receives and uses.

A Broadcaster provides the broadcast facilities to transmit the data from Data Providers to Data Subscribers.

A fourth player in some situations is a Service Bureau, an enterprise which handles administration of one or more components of the datacast sytem as a service to Providers, Broadcasters, and/or Subscribers.

As often happens, the same organization may play more than one role. For example, in applications where a head office is distributing information to branch offices, the same enterprise may be both Data Provider and Data Subscriber. In applications where a TV station is distributing audio clips to local radio stations, the TV station may be both Data Provider and Broadcaster.

A datacasting system needs to support all of these roles.

Data Flow. A very important aspect of datacasting is managing the flow of data from its source, where it is created or aggregated, to its destination, where it reaches the end user.

The data flow normally occurs in four stages:

- Retrieval of the data from its source
- Insertion of the data into the broadcast stream
- Extraction of the data from the broadcast stream
- Presentation of the data to the end user

These first two stages may or may not be closely coupled in time. For example: In a stock quote application new data may be retrieved from the source at prescribed intervals and immediately broadcast. In a real estate listing application new data may be retrieved from the source overnight and broadcast at intervals throughout the following day. In a tourist information application the data may be retrieved from the source at random times and then broadcast continuously.

Retrieval of the data from its source may be based on a “push” or a “pull” paradigm. The data may be pushed from the source whenever it is created or updated, or it may be pulled from the source at prescribed intervals.

Any individual data item may be inserted into the broadcast stream just once, or it may be broadcast some prescribed number of times over some prescribed period of time.

The second and third stages are inherently closely coupled in time. The data must be extracted at the time when it appears, which is essentially the same time as when it was inserted.

The third and fourth stages may or may not be closely coupled in time. In some cases the data is presented to the end user as soon as it arrives. In other cases it is stored and presented to the end user at some later prescribed time, or perhaps presented only when the end user asks for it.

A datacasting system must support the scheduling and implementation of these data flows.

Bandwidth Management. Depending on how the Broadcaster makes the trade-offs, the amount of bandwidth available for datacasting may be constant or may vary from time to time during the day. It may even vary on a second to second basis, depending on the detailed characteristics of the TV station's current programming, if the Broadcaster really makes an effort to maximize datacasting revenue.

The Broadcaster may choose to allocate this bandwidth among multiple Data Providers, possibly with different priorities. Each Data Provider may be allocated a certain amount of guaranteed bandwidth, plus a certain additional amount on a "best effort" basis. The "best effort" bandwidth may have a priority level associated with it (with the price charged to the Provider presumably depending on the priority level).

These allocations may depend on time of day and day of the week.

A datacasting system must support the allocations of bandwidth to different Data Providers, enforcement of the allocations, and in many cases metering and billing for the actual bandwidth used.

Security. End-to-end security is typically very important for enterprise-to-enterprise applications. In some cases it is simply a matter of assuring that no one gets the data unless they have paid for it. In other cases, such as distributing police reports and other government information, it may be essential to protect the data from disclosure to unauthorized persons for privacy reasons.

This means that it may be necessary to maintain security for the data at every step of the way:

- During retrieval from the source
- While in cache waiting for transmission
- During transmission in the broadcast stream
- While in the receiver waiting for use

In some cases the data must be kept confidential even from the Broadcaster. In other cases the Broadcast must be allowed to inspect the data to assure that it meets the Broadcaster's standards for the type of data it is willing to transmit.

A datacasting system must provide appropriate encryption/decryption and passcode mechanisms to assure the appropriate level of security.

Compression. Given that bandwidth is a valuable resource, in many applications it is worth while to compress the data before transmission and uncompress it upon receipt. This should typically be invisible to the end users.

A datacasting system must provide facilities for such compression and uncompression.

Error Prevention/Recovery. So far there has been relatively little experience with datacasting via terrestrial ATSC DTV broadcasts (with 8-VSB modulation), so the detailed error characteristics of such data broadcasts are not well known.

Other factors to take into account are that the impact of errors may vary widely from one application to another, and that the system needs to be resilient to the occasional failure of computers or other network components in the system, as well as to errors arising from corruption of the broadcast stream.

The situation is complicated by the fact that DTV broadcasts are a one-way communications channel. If a receiver wants to report errors and ask for a retransmission, it must do so through a separate back channel.

In general, some degree of error prevention and/or recovery is almost certain to be needed for certain applications. In some cases it may be adequate simply to rebroadcast each data item multiple times at prescribed intervals. In other cases a more elaborate forward error recovery scheme may be needed (above and beyond the forward error recovery scheme that is already a part of 8-VSB modulation).

A datacasting system needs to provide appropriate error prevention and recovery mechanisms.

Three-Component Architecture

It is clear from the discussion in the preceding section that a datacasting system must meet a number of requirements in order to properly support enterprise-to-enterprise data broadcasting. One key factor in the ability of the system to meet these diverse requirements is the underlying system architecture.

The ideal system architecture should mirror the problem space. In particular, it should explicitly recognize and support the roles of Data Provider, Broadcaster, and Data Subscriber.

An architecture for a datacasting system which meets this need very well is shown in Figure 1 below. It consists of three primary components:

- Data Source Server
- Data Hub Server
- Data Receiver

Data Source Server. The Data Source Server allows the Data Provider to specify the detailed scheduling for retrieval and broadcast of individual data items. This has advantages for both Data Provider and Broadcaster. The Data Provider can change the scheduling at any time without involving the Broadcaster in the process and going through a lot of red tape. The Broadcaster does not have to devote any clerical resources to this task. The specifications can be transferred automatically from the Data Source Server to the Data Hub Server for use by the Data Hub Server as needed.

The Data Source Server pushes data to the Data Hub Server as needed, and can optionally apply data compression, encryption, and forward error correction encoding if appropriate.

Data Hub Server. The Data Hub Server gets data items from the Data Source Servers and turns them over to the MPEG-2 Gateway according to the schedules obtained from the Data Source Servers.

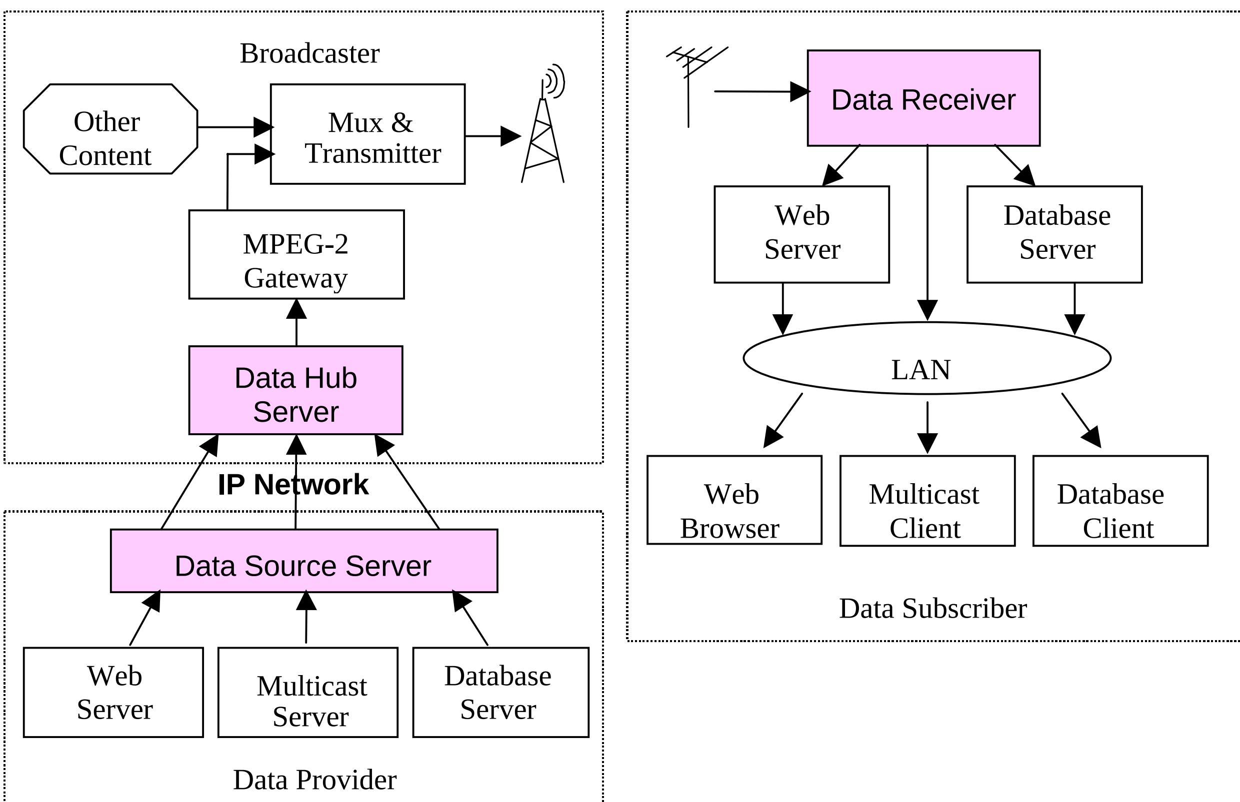

Figure 1. Three-Component Architecture for Datacasting

The MPEG-2 Gateway simply encodes the data into MPEG-2 transport packets and sends them to a multiplexor for actual insertion into the broadcast stream.

The Data Hub Server allows the Broadcaster to specify the bandwidth allocations and priorities of the various Data Providers, and it enforces these allocations. If the detailed schedule specified by a Data Provider causes an overrun of the Provider's allocated bandwidth, the Data Hub throttles it down to keep it within the allocation, and notifies the Provider that it has done so.

The Data Hub Server can apply compression, encryption, and forward error correction encoding if appropriate. It also meters bandwidth usage and can generate billing reports as needed.

Data Receiver. The Data Receiver extracts the data from the broadcast stream and applies decryption, uncompression, and forward error recovery as needed. It provides the Data Subscriber with a menu of the authorized items in the broadcast stream (generated from specifications provided by the Data Source Server and transmitted in the broadcast stream), and allows the Data Subscriber to select which ones to actually extract and which ones to ignore.

It then stores the items on local disk, retransmits them on a local LAN, or stores them in a local database, depending on the nature of the data.

Summary

Data broadcasting presents a potentially very valuable revenue opportunity for DTV broadcasters. Enterprise-to-enterprise applications are likely to be the best way in the near term to take advantage of this opportunity.

Implementation of such applications presents a number of challenges in such areas as data flow management, bandwidth allocation, security, data compression, and error recovery.

A three-component architecture is presented which supports the three primary roles of Data Provider, Broadcaster, and Data Subscriber. This architecture facilitates efficient, robust solutions to the system requirements.

References

[1] *ATSC Data Broadcast Standard* (Draft), ATSC Document T3-504, January 2000.

Printed from a paper delivered at the NAB Broadcasting Engineering Conference, Apr. 8-13, 2000.