



News On-Demand for Multimedia Networks

Gene Miller*, Greg Baber† and Mark Gilliland†

*NYNEX

White Plains, NY

†Dow Jones

New York, NY

ABSTRACT

A publisher of business and financial news and a regional telephone service provider are developing the means to produce, distribute, and retrieve multimedia news over wideband information networks. The goal is to deliver business news on-demand to corporate and financial services professionals in near-real-time using advanced network-based multimedia services. The new requirements of this application led to the creation of the following prototypes: a multimedia news programming format for live and stored information; a system concept for network-based multimedia information services; a flexible information structure to support composition, linkage, reuse, indexing, versioning, access control, and presentation control for live and stored multimedia; a groupware authoring system for audio-video capture, editing, indexing, and publishing; a digital information network with dynamically allocated stream and packet channels; a network-based multimedia store-and-forward service to provide notification, controlled distribution, access via authorizations, and database search and retrieval; a multimedia local area distribution system; and a user interface for browsing, search, notification and retrieval, with automatic hyper-linking to associated information and non-linear view/playback of audio-video material. The prototypes make efficient use of available multimedia and communications technologies, and are intended to exploit future network and media products as they become available. The resulting system serves as a high-performance testbed to study user requirements, to evaluate new technologies, to probe market needs, and to develop new multimedia service concepts. While the current work focuses on a specific application, it is expected to provide a model for developing a wide range of multimedia applications and services.

Key Words and Phrases: networked multimedia, news, digital video, information service, database, authoring, retrieval, wideband, messaging.

1. INTRODUCTION

In March, 1992 Dow Jones and NYNEX agreed to construct and operate a prototype multimedia news distribution service.

The goal of this collaboration was to define and prototype an end-to-end information service to deliver business news on-demand to corporate and financial services professionals. The service combines video, audio, text, and graphics in a format that permits browsing, searching, and user-notification. The systems, operating on a public wide-area network, deliver late breaking news with minimal delays. Users are able to quickly retrieve just the information they need at any time they want.

1.1 New requirements

A number of new requirements distinguish this application from conventional multimedia publishing efforts:

Time critical: To those people for whom the timeliness of information is critical, it is absolutely necessary that they be notified about an event in the few seconds after it has occurred.

Volume: Several hours of program material will be produced and distributed on a daily basis. Multi-gigabytes of highly compressed information will be accumulated every day.

Live and stored content: The application calls for live "broadcast" video to be presented along with material that has been digitized, compressed, and stored as a data file.

Prioritization: Content will be prioritized by urgency criteria established by the information provider.

Alerts: Users will be notified of important new information as soon as it becomes available.

Updates: Since news is inherently dynamic, numerous revisions of the same news story may be published within minutes of each other.

Aging: Certain news stories are discarded at the end of the day, while others are retained for months.

Archiving: Some information will be archived for permanent search and retrieval.

Dynamic user needs: Users have ever-changing information needs and need to be able to vary their "interest profiles" at any time. The information must be coded and categorized so that it can be retrieved at any time.

Distribution control: The information provider must be able to target each news story to a particular group of authorized subscribers, and be able to monitor the delivery.

1.2 Scope of work

The roles of the collaborating companies were defined at the beginning of the project. They decided to jointly develop the overall architecture, information structures, and service interfaces; Dow Jones would build and operate the multimedia authoring studio to produce and publish the news, provide live video news feeds, and build the subscriber site delivery system; and NYNEX would build and operate the prototype wide-area information network and the network based store-and-forward systems.

The following new concepts and systems were defined and prototyped in order to satisfy the requirements of this application:

- A new **programming format** for multimedia news on-demand (Section 3).

- An appropriate **system concept** for network-based multimedia information services (Section 4).
- A flexible **information model** to support composition, linkage, reuse, indexing, update control, access through authorizations, and presentation control for live and stored multimedia, including video, audio, text, graphics, quote data, and other media types (Section 5).
- A **groupware authoring** system to support rapid information capture, editing, indexing, and publication (Section 6).
- A **wide-area switched network** for multimedia information transport and session control, with dynamically allocated stream and packet channelizations (Section 7).
- A network-based **multimedia store-and-forward** service to provide notification, optimized distribution, and database search and retrieval. (Section 8).
- A multimedia **local area distribution** system and a new **user interface** for information browsing, search, notification and retrieval (Section 9).

To accommodate this prototyping effort, a new multimedia production studio and customer demonstration center was built by Dow Jones at their headquarters in New York City. These facilities are connected through wideband transmission lines to the storage, switching and distribution facility at NYNEX Science and Technology, in White Plains, NY.

Development began in March, 1992. Live and stored news stories were successfully distributed over a wide area network in November. Although certain operational changes are being made, the systems are expected to be deployed in a field trial during the summer, with customer evaluations continuing through December 1993.

2. BACKGROUND

As multimedia and technologies became available in the 1980's, researchers began to demonstrate their applications to news information services. Newspeek, [Brand 87] a personalized news information presentation service, was demonstrated at the MIT Media Lab. Dow Jones was a sponsor of this project. Later on this lab demonstrated high density presentation of real-time news (see [Bender 91] and [Masuishi 92]). NYNEX was a sponsor of these projects.

Publishers and broadcasters have recently announced trials and commercial offerings of several news oriented multimedia products, each with a near-real-time component. In 1990, Reuters Holdings PLC announced but subsequently canceled TV2000, a desktop teleconferencing service that featured call in interviews. [IDP 90a]. Also in 1990, Infotechnology announced and then canceled FNN:PRO, a system that delivered programming from FNN to quote terminals. [IDP 90b]. NBC announced that it will develop with IBM and NuMedia Corp. NBC Desktop News. This multimedia news delivery service will send customized clips from NBC news broadcasts to PC users. [Seybold 92].

2.1 Dow Jones

In 1990, the Dow Jones Information Services Group launched DowVision, an advanced information service that delivers the Dow Jones newswires, two press release wires, and the full-text of The Wall Street Journal directly to corporate desktops. The news is delivered in real-time and is tagged with MetaData, a term for the corporate ticker symbols, industry, government, product, subject, region, and market sector codes attached to each published story that allow users to filter and select only the information and news stories important to them. Applications created by independent companies called Alliance

Developers store the news stories at the customer site, forming a local database that is configurable, customizable, and instantly-accessible.

2.2 NYNEX

The work described in this paper began in the context of the 1991 US court ruling that permitted regional telephone companies to participate in the development and delivery of information services. In the 1992 Video Dialtone order, the FCC modified its rules to enable local telephone companies to participate in the video marketplace. The Dow Jones collaboration is part of NYNEX's overall program to promote new applications and services to meet customer needs, making use of emerging transmission and switching technologies. The multimedia information servers and gateways described here are expected to become key service elements.

3. DOW JONES PRODUCT OFFERING

In DowVision, news alerts in the form of relatively short textual stories are delivered to subscribers' desktops within seconds of their publication. For all but the most critical of applications, this can be considered to be news delivered in "real-time." For the Dow Jones/NYNEX multimedia project, we seek to duplicate the functionality of DowVision with news alerts encompassing all available media forms, including graphics, audio and video.

A system that accommodates all the additional media types gives editors much greater flexibility to choose the medium best suited to the news event. But the production of the content becomes significantly more complex and time-consuming, delaying publication of the news.

The result is a compromise. Reporters and editors take the extra time to produce a multimedia news alert with "acceptable" audio and video quality and sacrifice some level of timeliness. To compensate, the system definition was expanded to include facilities for distributing a "live" video stream, not unlike a wide-area multicasting video conference, for those events deemed too critical or time-sensitive to delay.

Therefore, normal priority news alerts are produced using commercially available desktop audio, video, and graphics packages and some custom authoring software. High priority news alerts or "hot" stories are distributed in textual form only, similar to the newswires systems, or in a live "televised" mode direct from the studio at Dow Jones.

4. SYSTEM CONCEPT

It is the authors' assertion that newswire systems like DowVision are essentially file transfer systems, that is, stories (files) created by editors on word processing systems are transmitted over wide-area-networks to computers at customer sites, where they are stored again as files, displayed on a terminal, or printed. A major part of the system described here is involved in the store and forward transfer of files.

The other major function of the system is the allocation of network bandwidth and system facilities for the transmission of isochronous (constant bit rate) streams of data for which there is no intermediate storage, and for which transmission resources must be allocated to permit synchrony between the creation and reception of the data stream.

System components include the Dow Jones Production Center, a wideband network, the NYNEX Media Service Center, and customer sites (see Figure 4.1). At the Production Center, media files are stored on one of two file servers. When a story is published, the file(s) are transferred to the multimedia storage

subsystem of the NYNEX Media Service. Then, they are forwarded to similar file servers located at the customer sites.

Stories can also originate as time-dependent streams of data from a video codec. When an appropriate command to the system is given, the output stream from the codec at the Production Center is passed through the network and delivered to the codecs at the customer sites.

4.1 Multimedia Information Content and Structure

The basic information unit of the system is a file. A file can contain data like ASCII characters, bitmap or structured graphics, a stock market quote or the results of a database search. Files can also contain time-based information such as digitized streams of audio or video, or a series of frames stored as an animation. Each file can contain only one media type.

Stories are collections of files. A simple story might contain only a single file, while a complex story might contain dozens of files with a varying mix of all media types. The decision for which types of media to use for each story is left to the editors.

Stories can also include streams of data from a video codec. It is important to distinguish the codec stream from the digitized video files. Files can be delivered over a packet network where the delivery speed or time is not guaranteed, but codec streams cannot. If there is insufficient bandwidth in the network to deliver the stream, the codec at the customer site will become data starved and the video will appear broken up and erratic, or not at all. A file, on the other hand, will simply arrive later than expected, but can be played back without distortion.

Control information for both files and streams is stored in a data structure called an Object Control Structure (OCS) (see Section 5). For files, the OCS contains a "pointer" (full path name) to the data on the file server. For the codec stream, the OCS contains a pointer to the network channel on which the stream will appear, the time the stream will start, and its approximate duration.

OCSs can be combined in varying hierarchies (i.e. stories can contain other stories) to permit the most efficient packaging and reuse of the information content as well as the network and system resources. It is this ability to combine both stored and

isochronous data in a structure that permits such a rich blending of these two dramatically different media forms that makes this system unique.

4.2 Dow Jones Production Center

At the Production Center, the "type" of media stored in the file determines the server on which the file is stored. Digitized audio and video files, as well as other time-dependent media types, are stored on the audio/video file server. Besides having storage capacity to accommodate these files, which can be as large as 1.3 Gigabytes for a one hour video, this server can deliver isochronous streams over an Ethernet local-area-network at rates as high as five or six megabits per second to several users simultaneously. Clients of this server can therefore store and play back these files in real time.

Text, graphics, and other non-time-based files are stored on a Novell file server. Segregating time-based and non-time-based files in this fashion permits each server to be tuned with respect to block size, buffer capacity, etc., for that class of media.

OCSs, limited by definition to 64 kilobytes, are stored in a network-model database under the control of a system called News Publisher. News Publisher is a custom application designed by Dow Jones to control the creation, modification and deletion of OCSs and files within the Production Center, and to permit the publication of these files to the NYNEX Media Service Center.

4.3 Transport Facilities

A wide-area network connects the Production Center and customer sites to the NYNEX Media Service Center. The current implementation of this network is over private line T1 circuits, but work is under way to change to a switched wideband service implemented on Primary-rate ISDN (see Section 7).

4.4 NYNEX Media Service Center

The NYNEX Media Service Center provides centralized storage as well as communications and information management for the system. These services are designed to be generally applicable to a broad range of information including the Dow Jones product offering.




Figure 4.1: Overall Architecture

The Transport and Session Server controls and monitors communications paths used by the system (see Section 7).

The NYNEX Media Server (see Section 8) provides central database, notification, transaction, and distribution services. It controls the transfer of media files into the multimedia storage subsystem, and allows storage, retrieval, searching, modification and deletion of the OCSs and media files.

4.5 Dow Jones Customer Site

Dow Jones operates the customer site equipment, which is essentially a mirror image of the file servers and communication equipment used at the Production Center, except that the custom application called News Manager is present to coordinate end user requests and notifications. This will be covered in more detail in Section 9.

5. INFORMATION MODEL & STRUCTURE

A key objective of the system is to provide a common information model that would support the various kinds of live and stored data produced by Dow Jones multimedia applications while still providing uniform access and control of the data at the NYNEX Media Server. In addition, the model should not only be adaptable to Dow Jones's changing requirements, but also fit the needs of other information providers. Looking ahead, this same model should also be suited for the wide range of new applications that are being investigated in our labs and elsewhere (e.g., distance learning, collaborative design, video catalog shopping, etc.). Thus, flexibility and extensibility of the data model are important. Other characteristics that should be supported by the information model include efficient transmission, information reuse, indexing for fast search and retrieval, update control, privacy and security, and presentation control for multimedia.

5.1 The Media Object

The main entity in the information model is the Media Object, an abstract entity that encapsulates numerous data formats and mediums and is uniformly accessible by upper layers of the system. A Media Object contains two kinds of information: 1) an Object Control Structure (OCS); 2) an optional media data set (file or stream) which is the raw data payload of the Media Object; A Media Object may optionally point to other media objects to provide object composition. We defer the discussion of composite pointers to section 5.2. (see sections 5.4 and 5.5 for detailed description of Media data sets and OCSs, respectively).

An important characteristic of this model is that media files and control information are kept separately. The physical as well as the conceptual separation between control & summary data, and the actual "product" data — bears several advantages: it enables the light-weight OCSs to reach their destination quickly without the need to wait for the large files to arrive. (OCSs will always be less than 64,000 bytes, while a 30 minute video clip could be more than 320 million bytes.) Moreover, it enables OCSs to be distributed without their associated files. This is particularly important for customer sites where a large number of OCSs would be available for browsing and searching, and where a substantially smaller number of files would be fetched on demand for access from the local cache. (See section 5.7 for consistency implications of this feature). The separation also allows OCSs to be managed in the NYNEX Media Server using standard DBMS technology (see [Thur 92] for discussions of Object Oriented DBMS) and keep the large media files externally on a suitably efficient storage system that would support real time stream playing and other features that are unique to multimedia data sets (see [Rangan 92] and [Rangan 93]. Finally, the multi-

media applications that write and read the media files do not have to recognize media objects, since the system hands them the raw media files or streams. This allows any off-the-shelf multimedia applications to be integrated into the system. This is in contrast to the approach taken by some commercial database vendors where "multimedia" data must be accessed through proprietary interfaces.

One problem with this approach is that the various servers (Media Server, News Manager, and News Publisher) must maintain the OCS and its associated file consistent with each other, and ensure that an OCS can always reach the most recent version of its associated file, even if it is currently being transferred, or does not reside in the same node. The problem of object consistency is discussed in detail in section 5.7.

5.2 Composite Objects

In order to support the creation and representation of compound stories, composed of multiple possibly independent components that can be manipulated separately and perhaps added incrementally — the information model must support object composition. Moreover, a single object may be used by multiple different composite objects. For example, a video segment can be used by multiple different news stories. This requires the ability to share objects among multiple composite objects.

Thus, our model supports both object composition and shareability of sub-objects and can be described as a DAG (directed acyclic graph) with two kinds of nodes: a Simple Media Object, which has no outgoing edges, and contains a simple data set, which is classified as either stored media file (FILE) or as a live media stream (LIVE); and a BINDER object, which has at least one outgoing edge to other objects, of either BINDER, FILE, or LIVE type.

Figure 5.1 shows several possibilities: A picture FILE object (1) is contained in no binders; a text FILE object (2) is contained in one BINDER object (6), which is itself contained in another BINDER object (8); a LIVE video stream object (3) is contained in the same BINDER object (6). A video FILE object (4) is contained in two BINDER objects (8 and 7); and a audio FILE object (5) is contained in BINDER object (7).

For the Dow Jones News application, a BINDER object represents a story or program set, a FILE object represents a component media file of the story (e.g. a compressed digital video segment, or a text or graphics file), and a LIVE object represents a live video feed that is distributed to one or more sites.




Figure 5.1: Simple and composite media objects. The numbered circles are the OCSs.

Media Objects may be stored, retrieved, replaced, or deleted from the NYNEX Media Object Server. Applying any of these operations to a Binder may cause the operation to be recursively applied to all the objects contained in the Binder. This is discussed in length in section 8.

5.3 Object Versions

Versioning of objects is desirable for any design environment, where authors may need to revise several times their artifacts, or reuse pieces for creation of a new artifact. In addition, versioning can reduce the amount of space requirements. Within the context of the Dow Jones news application, versioning is imperative, since late breaking news is frequently updated, as new information is obtained, or as MetaData codes are added.

In our model, each object has a special attribute called modification (mod) number, which is maintained by the system and keeps track of multiple versions of an object. Object modifications are considered non-destructive. A later mod can safely replace an earlier mod. For example, the event list (see below) in a later mod is a normally a superset of the event list in the earlier mod. This assures that the binder object can safely reuse child objects.

5.4 Multimedia Data Sets

The raw materials of a news story are data sets, each containing data of a specific format. The data set can be a stored file or a continuous stream that occurs on a channel during a particular time interval. A data set may contain raw device control data, e.g., pixel values or sound samples, as well as auxiliary format and structural information. Data sets may be compressed or uncompressed.

Data sets are usually associated with existing data handling software or hardware systems. Examples include: An ASCII text file; A Macintosh PICT file; An AVS compressed video file; A Px64 compressed video stream.

Each data set belongs to a general media type, (e.g. VIDEO) and a specific media type (e.g., DVI), and has specific attributes (e.g., rows, columns, frames per second, bit rate).

5.5 Object Control Structure

The OCS for a media object is the collection of attributes for that object. OCSs contain both the system attributes as well as the publisher attributes. System attributes are mandatory and are contained in each OCS. Publisher attributes are specifically defined to meet the needs of each information provider.

System attributes are used for access control, version management, storage management and distribution management. The Application Programmer's Interface document [NYNEX 92] contains a detailed list of system attributes in the OCS.

The set of attributes determined by the publisher is extensible, i.e., new attributes may be added at any time, although full schema evolution is not supported yet, so deletion and modifications of attributes and of type definitions is not supported. To enable uniform extensibility, several data types (e.g. Date, String, Blob) are available for defining new attributes. For each data type, there is a set of comparison operators and a standard undefined value.

5.6 Event Lists and Multiple Views

A single object can be accessed as a set of logical sub-component. For example, a video file may be defined to have 5 logical "scenes", where each scene is tagged with a start and end frame number. This information is currently encoded in an object's event list. With this representation, it is possible for a

story to "play" scenes out of order (5,1,3,2,4). Various stories may access different scenes in the same object, thus implementing multiple views of the same underlying media objects. In addition, a parent object may inherit the play list of all of its children.

5.7 Object Consistency

Having defined an object model, we have to define a set of properties that define the consistency model. Once defined, the transaction management layer (discussed in section 8) has to ensure that the system is always in a consistent state.

The following properties must all hold invariably in NYNEX's Media Object Server and Media Data Server:

1. A single file must always exist in its entirety or not exist at all in the Data Server. Likewise, a single OCS must always exist in its entirety or not exist at all in the Object Server.
2. For a FILE object, the OCS and its associated file must both exist in the Object Server and Data Server, respectively.
3. A binder must have all of its child objects present at any time (recursively).
4. The modification (mod) number of a child object must be greater or equal to the mod number referred to in the object list of its binder.

Note that any object (including binder objects) may reside in the Media Server without its parent binder, even if it originally had one. It may, or may not, be joined at a later point by its original or by new parent(s). The reason for this relaxation is to enable more efficient consistency management. Section 8 elaborates on this issue.

6. DOW JONES PRODUCTION CENTER

The Dow Jones Production Center is the source of all of the news material for the system. It consists of 7 components: four Production Stations (a Capture station, an Editing station, a Text/Graphics station, and an Authoring station), an Audio-Video File Server, a Text/Graphics file server, and a specialized application called News Publisher, which is the repository for all control information in the Production Center. All of the stations are 50 MHz Intel 80486 PCs and run on an Ethernet local area network. The Production Stations are equipped with Intel ActionMedia II DVI capture and playback cards.

The **Capture** station digitizes and compresses the analog audio and video and stores the "raw" media files on the Audio-Video File Server in AVSS (Intel's Audio/Video Support System) format. Video is compressed in RTV 2.1 format. Audio and video can be input to the system from any live source (satellite, cable, etc.) or from standard analog tape.

The **Editing** station is an off-the-shelf DVI-compatible Digital Video Editing (DVE) system used to post-produce the raw audio and video files created by the Capture station.

The **Text/Graphics** station is a standalone editing station used to create or import all of the text- and graphics-based media files for the system. Ultimately, it will be replaced by a network connection to the text and graphics systems used in the production of The Wall Street Journal. Both the text package and the graphics package are off-the-shelf.

The **Authoring** station is a custom software application that allows an editor to create a multimedia story by reviewing and selecting the text-, graphics-, audio-, and video-based media files on the file servers. The files are combined under display formats or templates, a headline is written, and the MetaData and authorization codes are attached, and a command is sent to News Publisher to begin transferring the story to the NYNEX Media Server.

The **Audio-Video File Server** is an off-the-shelf server that combines special hardware and software to provide guaranteed delivery of isochronous streams of audio and video to and from the Production Stations. The server and client software was developed by Starlight Networks Inc.

The **Text/Graphics File Server** runs standard Novell NetWare v3.11, and is used to store all of the text and graphics files.

News Publisher is the station that controls all of the activities in the Production Center. All of the Production Stations communicate with News Publisher to negotiate the creation, deletion, and modification of Object Control Structures for the media files, and to interact with the NYNEX Media Service. News Publisher stores all of the OCSs in a small local database, and assigns all system-defined values in the OCS.

6.1 Story Flow

Like other established editing environments, the production of multimedia news stories is, at first glance, hierarchical. A story passes through several levels of editing until it reaches the most senior level editors or their deputies, who have final responsibility for publication of the article. The story appears to flow in linear fashion from one person to the next.

However, a more detailed analysis of the traffic flow shows it to be quite serpentine. Editors often send stories back to reporters for minor changes or wholesale revision. Two or more reporters collaborate on an article, frequently sharing incomplete drafts. A reporter checks the figures on a chart or graph to compare them against the text in an article. Senior editors review work-in-progress at all levels.

To accommodate for these traffic flows, the Production Center was designed as a groupware collaborative editing environment, under control of the News Publisher application.

Once a media file has been edited and is ready for inclusion in a story, the operator of the station requests the creation of an Object Control Structure to News Publisher. News Publisher

creates the OCS, stores it in the local database, and passes an object identifier (OID) back to the requesting application. The OID and all relevant information about the article (headline or caption, date, revision, etc.) are made available to all Production Stations. Applications can then retrieve the caption for that media file and display it to their users.

6.2 News Publisher API

All control of the OCSs is mediated by News Publisher. We have defined a News Publisher Application Programming Interface (NP-API) through which each Production Station interacts with News Publisher and the NYNEX Media Server. The NP-API includes functions which allow Production Stations to search for media files with specific attributes both locally at the Production Center or remotely at the NYNEX Media Server.

6.3 Modification

For each OCS, News Publisher maintains a count of the number of times the OCS has been modified since last published. This is known as the operational modification (op-mod) number, and is different from the mod number used in the transfer of the media files to the NYNEX Media Server.

From the perspective of the Production Center, the mod number indicates the number of times that an OCS and its related media file have been published to the NYNEX Media Server. Conversely, the op-mod number is used to indicate the number of times that the OCS is modified between publications.

The Production Station applications compare op-mod numbers for each OCS to determine whether they have been modified since the last poll.

6.4 The Authoring Station Application.

The Authoring application is key to the production process, and was custom designed by Dow Jones specifically for this project. The Authoring station displays lists of all of the media files and their OCSs available on the file servers, and allows an editor to combine multiple media files into a multimedia story.

The Authoring station displays a story as a two-dimensional matrix of media objects, organized similar to a spreadsheet. Each row of the matrix represents a media type, such as text, graphics, video, etc. Each column of the matrix represents a "scene" in the story. In this way, editors can create a story with multiple scenes, such as a video segment followed by a graphic with an audio track, followed by another video clip and a piece of text.

Also, the Authoring station uses custom viewers to preview and select objects or sub-scenes in video objects.

Most importantly, editors can "tag" each event with different MetaData codes. A multimedia story about the auto industry may contain different segments about GM, Ford, Chrysler, etc., with each segment coded specifically for that company and the story as a whole coded as an auto industry story.

7. MEDIA TRANSPORT and SESSION SERVER

The NYNEX Media Transport and Session Server provides application independent transmission management and session services. The underlying transport is provided by private line services (T1) or switched wideband service, such as primary rate ISDN.

The Transport Manager layer allocates logical channels within the transport medium to allow media with differing performance requirements to be transmitted appropriately. Quality of service characteristics include bandwidth, error rate, and latency. The basic function used for the current application is the dynamic allocation of isochronous stream and packet channels (refer to

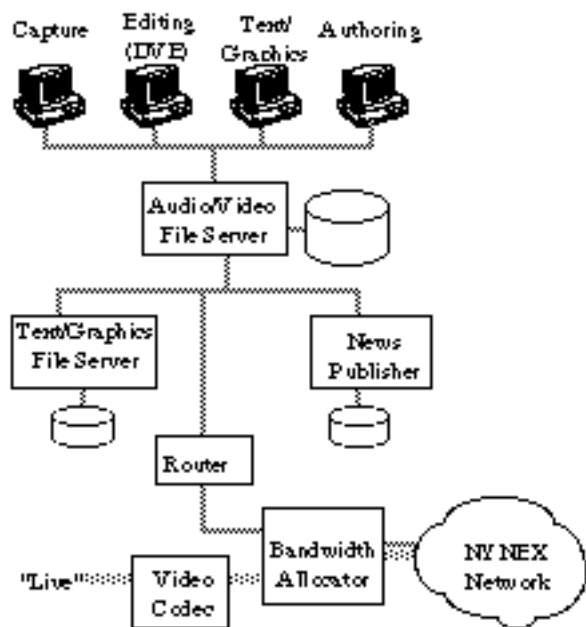


Figure 6.1: The Dow Jones Production Center, located at the World Financial Center in New York City, is the source for all news content in the system.

figure 4.1) The Session Server layer prioritizes data transfers and manages connections.

A key component is an inverse multiplexer, called the Bandwidth Allocator, that can dynamically call up multiples of D0 circuits and aggregate them into higher bandwidth channels for such functions as signaling, rapid file transfer, and isochronous video streams. Multiplexers are at each end of the communications links, and operate under the control of the transport and session server.

The packet port of each Bandwidth Allocator is connected to a TCP/IP packet router at each location. Stream video ports at each Dow Jones site are connected to video codec, and at the NYNEX center they are connected to a video stream distributor.

8. NYNEX MEDIA SERVER

The NYNEX Media Server is responsible for the storage and distribution of all media object, including the OCSs, and the underlying FILE and LIVE multimedia data sets.

The NYNEX Media Server consists of three server modules: Transaction Server, Object Server and Data Server.

8.1 Transaction Server (TS)

This is a central component of the Media Server—to properly and efficiently distribute objects from the publishing site(s) to the Media Server to the multiple customer sites that are authorized to receive the information. In addition, it provides scheduling services including queued task management, prioritization, and coordinated information transfer. Finally, it interfaces with the Object Server for information access, and provides resource management.

In order to ensure the consistency of the objectbase (as defined in section 5.7) all operations involving media objects are executed within a context of transactions. We first describe the set of services provided in this component, then proceed with presentation of a transaction model specifically designed to support operations that deal with distribution of media objects and the specific recovery procedures, followed by the scheduling services provided.

When the *MsPutObject* primitive request is issued at the publishing site (News Publisher), the distribution process is initiated. It consists of a single upload operation, and a set of download operations, one per customer site, determined by the authorization group field of the object.

Any transfer of media objects (either upload or download) is a two-step process: First, the OCSs of an object are transferred, followed by transfer of its payload data set. In transferring the OCSs, composite objects are traversed in breadth first fashion, starting at the top-level BINDER object. Note that if the OCS is already present (with the required modification number), then the transfer is not necessary.

Once the OCS transfers are complete, the necessary data set transfers are initiated. In an upload operation, each new media file is transferred. For download operation, transfers are done only for those files which have their optional flags set in the object list of the parent, or in the *MsPutObject* request itself. These option flags include: 1) Specify that only the upload is to be done. The downloads will be deferred 2) Wait for user notification before doing a download.

When an *MsDeleteObject* primitive request is issued at the publishing site, a delete indication is sent to all authorized customer sites. The delete transaction is completed only when all sites have acknowledged the delete indication that is generated, or

when a time-out has occurred. The actual deletion of objects from the central server is handled by the Object Server, discussed below.

Transaction Model: Like traditional transaction systems, A transaction is defined as a set of operations that transfer the objectbase from one consistent state to another consistent state, both under normal operation and in the event of failure or concurrent operations. However, unlike standard transaction processing systems, a typical transaction in the Media Service may take very long (e.g., a transfer of simple video object from the Publishing site to the NYNEX Media Server may take one hour); it may involve large amounts of data (e.g., several hundred megabytes of compressed video data); and may involve highly-structured composite objects. Therefore, we deviate here from the standard transaction model [Kim 90] and use a more suitable long transaction model [Dayal 90].

First, transactions are not atomic (“all or nothing”). This is because the distribution of a composite binder with large number of sub-objects may take hours to complete, and it is both unacceptable and not desirable to rollback all the sub-objects that were completely stored due to an abort that occurred while accessing a particular sub-object. In order to still obey the consistency constraints laid out in section 5.7, we allow for sub-objects that were transferred completely to remain in the objectbase even if their parents or any of their “sibling” objects did not succeed.

The second deviation from the standard model is relaxation of the isolation property. Due to the long duration nature of the transactions, we allow for multiple ongoing transactions to perform access operations on the same (composite) objects in an interleaved manner so that overall throughput can be improved. Therefore we allow this to happen, as long as the partial results are guaranteed to not rollback in a failure event. This means that simple objects can be accessed even when their parent binders have not been transferred yet (assuming those objects have been transferred themselves completely, following the first consistency constraint). However, binder objects should not be accessible to other transactions until the original transaction has ended to avoid the “cascading aborts” phenomenon.




Figure 8.1: Roll back of distribution

Finally, since some transactions span multiple hosts, they can be regarded as distributed transactions. Therefore, a two phase commit occurs.

Rollback of Distribution: Figure 8.1 illustrates what happens when a transaction aborts in the middle of an upload. It is assumed that the upload of file 3 aborts due to time-out, network failure, publisher request, lost file, etc.. Only consistent objects are committed (shown as solid black). The inconsistent objects (white) are rolled back to their previous state. If the rolled-back objects are new, the OCS and associated file are deleted. If the rolled-back objects are just updated (new mod number), then the OCS prior to the update is restored, but the file is untouched. This rollback strategy is followed at the NYNEX Media Server as well as at all the sites of the download.

Scheduling: Due to the long duration of some operations, and due to the variance in the “degree of freshness” that stories should have (e.g., a late breaking story, vs. a feature story), the Media Server has to provide a flexible, priority-based, scheduling mechanism that supports a wide range of options, some of which are user-initiated. Scheduling options which are supported by the Media Server include:

8.2 Object Server (OS)

OS provides persistent and consistent storage and retrieval of OCS information, and maintains the consistency between the OCS and the files in the Data Manager. In addition to stored data, OS manages OCSs of LIVE objects. OS is called from the Transaction and Distribution server and calls the Data Server and the Video Stream Distributor

The Object_Add service primitive is normally called from an **MsPutObject** request initiated at a publishing site, handles a single object addition. Thus, the addition of a composite object is a multi-step process managed by the transaction and distribution server. OS guarantees the atomicity of a single object addition, and assures the consistency of the OCS and the corresponding media file.

The Object_Delete service primitive marks an object as available for deletion. Deletes are not necessarily recursive, that is, deleting a binder does not necessarily purge its contents. An expiration mechanism is normally relied upon to garbage-collect orphan objects. The Object Server maintains an expiration date on all objects. When an object is expired, is automatically deleted, but only if it is not contained in any binder.

The Object_Query service primitive is an important aspect of the system, as it largely determines the usability of the system for archival information browsing retrieval. For each type of attribute in the OCS, there is a set of defined predicates. These include:

- For all **number** attributes (expiration date, publication date, modification date, etc.), the usual **order** predicates are valid: =, !=, <, <=, >, and >=.
- For all **text** attributes (abstract, caption, by-line, etc.), a **substring** predicate is defined, where matching is case-insensitive, and words are delimited by certain white space and punctuation.
- For all **key set** attributes (company ticker symbol, government codes, etc.) a key match predicate is defined.

When a predicate is applied to an OCS in which the attribute is undefined, then the result is false. Predicates can be combined with the Boolean operations AND, NOT, and OR; and order of application may be controlled with parentheses. An Object_Query expression formed in this way is applied to the set of all committed OCSs, and the resulting matches are normally return in reverse chronological order.

Schema Evolution: In order for the system to be extensible, there has to be a mechanism that allows the information provider

to dynamically change the structure of objects, including pre-existing objects [Kim 90]. For example, addition of new search categories. Schema evolution is currently supported to a limited extent, in that only additions of new attributes are supported, and only publisher attributes can be modified.

Live Stream Management: An OCS for a LIVE object is normally distributed prior to the occurrence of the LIVE transmission to notify all sites and their users of the LIVE event. The OCS of a LIVE object contains the expected start and stop times of the LIVE stream.

Just prior to when the LIVE program stream is to start, a **Live_On** alert is sent by the Publishing node with the access authorization codes set to the affected sites. Likewise, when the LIVE stream is to stop, a **Live_Off** alert is sent by the Publishing node.

These alerts cause the Session Manager and Session Agents to allocate the necessary isochronous channel bandwidth for the video stream.

This alert also allows the News Manager applications at each site to distribute the video to interested users.

8.3 Data Server (DS)

DS is responsible for managing, storing, and retrieving Multimedia data sets of type FILE and LIVE. DS is shielded from knowledge of the object model, and deals only with raw multimedia data. DS interfaces only with the Object Server.

9. CUSTOMER SITE

A customer of the Dow Jones multimedia service will be provided access to multimedia information (as published by the Production Center) and to DowVision (Dow Jones’ real time text-based integrated newswire feed). Although the two services utilize vastly different delivery and transaction mechanisms, and are currently delivered on logically separate feeds, the two information streams share several important characteristics: every story event contains publication date (for aging control), modification information (for database update synchronization), MetaData codes (for content indexing), and headline captions (posted as an event notification to “interested” users). Specific details of the DowVision service are outside of the scope of this paper.

9.1 Delivery System

The customer site delivery system, Multimedia News Manager, consists of 5 components shown in Figure 9.1: DowVision Comms Server, a specialized application called News Manager, an Audio-Video File Server, a Text/Graphics file server, and Client Workstations.

DowVision Comms Server provides X.25 feed termination.

News Manager is the station that controls all of the activities in the Customer Site. The News Manager is a client of the NYNEX Media Service and receives feeds from the X.25 Comms Server. News Manager communicates with all of the Client Workstations to send alerts, and provide headline feeds. News Manager stores all of the OCSs in a local database, and managers files stored in the Audio/Video File Server and the Text/Graphics File Server. It also stores all of the DowVision data in a local database.

The **Audio/Video File Server** is a Starlight Networks media server to provide guaranteed delivery of multiple isochronous streams of audio and video to the Production Stations. In order to satisfy real-time constraints on each of several Ethernet segments, it also serves as a bridge between these segments.

The **Text/Graphics File Server** runs standard Novell NetWare v3.11. It is used to store all Text/Graphics files.

The **Client Workstations** include digital video/audio hardware, run special retrieval software and communicate with all other components at the customer site.

Additional components include **Router, Codec, and Bandwidth Allocator**.

Session and transaction control for the customer site are logically similar to those described for the Production Center. However, a News Manager client site has access to a limited subset of NYNEX Media Server functions, specifically excluding the ability to create, modify, or delete objects, or to read or change the authorizations of clients.

If authorized for reception, the News Manager will be notified by the NYNEX Media Server of any OCS publication (which may consist of either a Simple Media Object or a complete BINDER hierarchy). The News Manager will download and examine the indicated OCS in order to determine exactly what actions must be undertaken. (As described elsewhere: only an OCS with a newer revision than that locally held need be downloaded. Furthermore, transaction control flags in a BINDER's Object List indicate whether or not constituent FILES need be downloaded immediately or should pend on a specific user request.) Upon successful download, or upon error, the News Manager responds with appropriate completion status. It must then ensure the consistency of the local database. Results of this site's download are reported by the NYNEX Media Server to the Publication Center as part of the overall publication transaction status matrix.

The News Manager application provides broadcast alerts (headline captions with metacode) notifying all client stations when a multimedia story has been downloaded and is available for viewing. At the client workstation, the user application filters these broadcast alerts based upon attached MetaData codes and local user-defined preference profiles. Depending upon this metacode analysis, the alert may be routed to one or more view

windows, or it may simply be discarded. (Note that DowVision newswire text is handled in a similar manner.)

Additionally, multimedia story data is always accessible by means of database search. Any user may request a list of previously published stories available at his/her local News Manager database. Likewise, the NYNEX Media Server can be searched (limited to stories for which the particular customer site has been authorized) allowing access to material which has previously "aged" out of the local News Manager database. Search criteria may include metacodes, keywords (from headline alerts), media type (of embedded FILES), and publication date (or date of last update); such criteria can be expressed using complex Boolean operators.

9.2 User Interface

The user interface of the client workstations is responsible for providing a seamless integration of the various news event sources and underlying data types. It is built upon a News Manager API (NM-API) which provides access to the various sub-systems within the local Multimedia News Manager server (e.g.: OCS database, digital media stream server, and broadcast story headline event stream). Finally, it should be noted that the NM-API is intended to allow Dow Jones Alliance Developers to provide alternative third party interfaces and unique applications for access to the multimedia data service.

The client workstation user application has been designed to provide a customizable interface, as determined by individual preference. Configuration of User Profiles (i.e.: metacode filters of arbitrary Boolean complexity) is possible. However, in addition, each customer system is provided a set of "default" Dow Jones defined profiles that represent various pre-packaged service offerings such as Equities, Foreign Exchange, Capitol Market Reports, etc.. Each profile filter is assigned a unique "view window" for presenting temporally ordered story headline caption events as received.

There are several classes of profile View Window types including: DowVision style textual headlines list and preformatted tableaux. Color Plate 1 illustrates a tableaux view containing icons, headline captions, and button-triggered viewer options. Color Plate 2 illustrates the simple "waterfall" Headline view.

A point and click method allows the user to select any specific story event for further detailed examination. When selected, a story may be presented as its own tableau (allowing access to and selection of underlying sub-stories and/or constituent media files). See Color Plate 3. Or a simple story may directly invoke appropriate media viewers. See Color Plate 4. Standard Media viewers include support for: digital video, digital audio, graphics, text, and the LIVE video channel.

Where appropriate, the viewers for time-based media allow for standard VCR-like actions (e.g.: play, pause, rewind, mute), as well as direct non-linear access to key event frames. This indexed play feature is provided by means of a Program Guide attached by the Production Center to every story (derived during authoring process from scene lists embedded within the underlying media file OCSs). See Color Plate 4.

10. CONCLUSION

The objective of this project was to deliver timely news in response to end-users needs. As work on this project progressed, many specific requirements were defined and were subsequently met. The project has identified the system elements that are needed for an advanced news information service, and it has shown that such a system can be built with existing technologies.

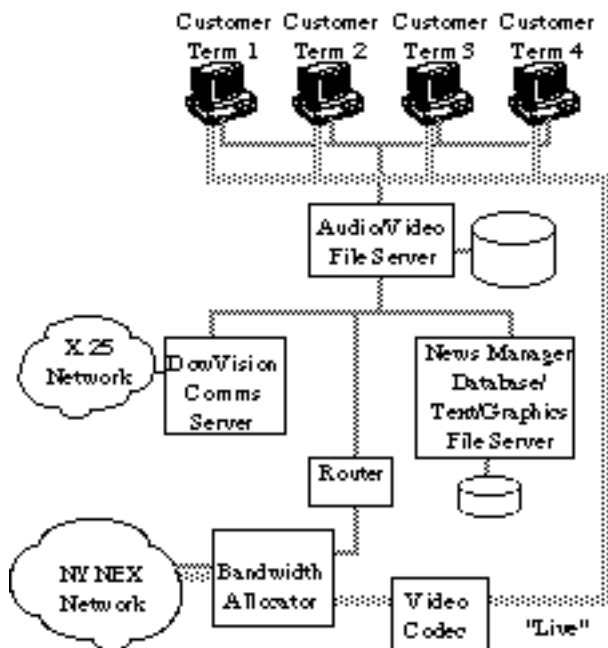


Figure 9.1: A typical customer configuration

As stated in the introduction, live and stored news stories are being sent to test sites over a wide area network. At the time of this writing, all key functions of the system have been implemented. The system will remain on-line until July, 1993 to permit a wide range of news story formats to be tested. Certain operational changes are being made to improve system performance and to facilitate maintenance and failure recovery. The systems are expected to be deployed in a field trial during the summer, with customer evaluations continuing through December 1993.

Scalability, reliability, and security requirements are being examined in anticipation of a future service offering targeted for a large number of customers sites. This will require a fault-tolerant distributed server architecture for the NYNEX Media Server. Future enhancements of the Server may also include: direct capture of live streams; direct play of audio-video files with motion control; direct modification of audio-video files; support for stand-alone customer terminals (e.g., TV monitors); full schema evolution; support for real-time collaborative editing; and support for real-time remote presentations.

In addition, Dow Jones has identified several areas for improving both production and consumption of the content.. Specifically, enhancements to the Production Center and Customer Site software will include a visual-based authoring environment (drag-and-drop), the inclusion of real-time market data and other financial instruments, a client-side API for third-party applications other than Dow Jones, and support for distributed Production Centers.

Both companies are seeking to apply this work toward additional multimedia initiatives. Participation in industry forums and standards groups is viewed as essential to the continued development of these concepts.

REFERENCES

- [Bender 91] W. Bender, H. Lie, J. Orwent, L. Teodosio, and N. Abramson, "Newspace: Mass Media and Personal Computing", *USENIX Conference Proceedings*, Nashville, TN, June 1991.
- [Brand 87] S. Brand, *The Media Lab: Inventing the Future at MIT*. Viking Penguin, N.Y., 1987.
- [Dayal 90] U. Dayal, , M. Hsu and R. Ladin, "Organizing Long-Running Activities with Triggers and Transactions," *Proceedings of the 1990 International Conference on Management of Data*, ACM, 1990.
- [Dow Jones 92] "Multivision Product, Preliminary Product Specification", Dow Jones, April, 1992.
- [IDP 90a] IDP Report; Copyright Knowledge Industry Publications Inc., June 22, 1990.
- [IDP 90b] IDP Report; Copyright Knowledge Industry Publications Inc., June 22, 1990.
- [Kim 90] W. Kim, "Research Directions in Object-Oriented Database Systems," *Proceedings of the 9th ACM Symposium on Principle of Database Systems*, ACM, 1990.
- [Masuishi 92] T. Masuishi, D. Small, and R. MacNeil, "6Kx2K Display Prototype", *SPIE Conference 1664 - High Resolution Displays and Projection Systems*, San Jose, CA, Feb. 1992
- [NYNEX 92] "Media Server, Application Programmer's Interface", Ver 1.18, NYNEX Science & Technology, Oct. 1992.
- [Rangan 92] P. Venkat Rangan, Harrick M. Vin, and Srinivas Ramanathan, "Designing an On-Demand Multimedia Service", *IEEE Communications Magazine*, Vol 30, No. 7, July 1992. Pages 56-64
- [Rangan 93] Harrick M. Vin and P. Venkat Rangan, "Designing a Multi-User HDTV Storage Server", To appear in *IEEE Journal on Selected Areas in Communications*, January 1993
- [Seybold 92] NBC tries 'Desktop News', *Digital Media, a Seybold Report*, Dec. 14, 1992
- [Singh 91] A. Singh and R. Silvers, "The NYNEX Shuttle", *IEEE 1991 Conference Proceedings*.
- [Thur 92] B. Thuraisingham, "On Developing Multimedia Database Management Systems Using the Object-Oriented Approach", *Multimedia Review*, Vol. 3, No. 2, Summer 1992.

News On-Demand for Multimedia Networks
by
Gene Miller, Greg Baber & Mark Gilliland

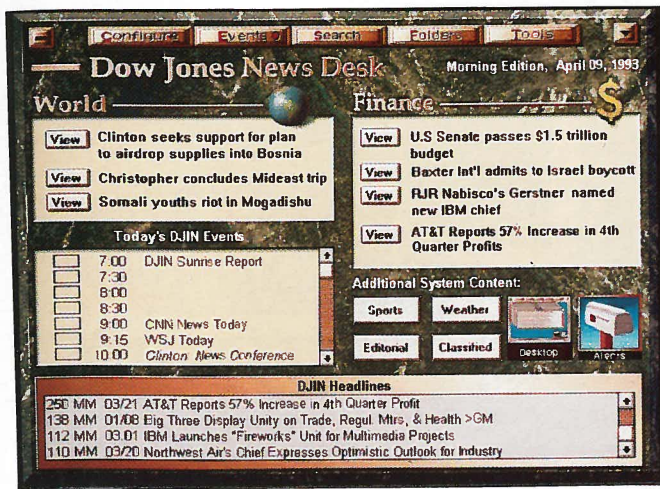


Plate 1

The *News Desk* is a presentation tableau designed to provide a unique electronic publication interface. Multimedia stories are automatically grouped by general topic (e.g.: World News) or service provider (e.g.: ESPN Sports). Additionally, a Daily Events panel allows alarms to be set for scheduled Live broadcasts.

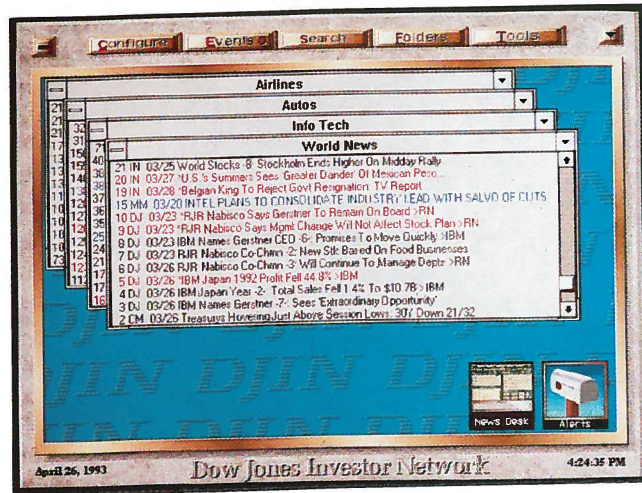


Plate 2

The waterfall Scrolling Headline News Panels provide an alternate interface which integrates both Multimedia stories and DowVision text stories. Each News Panel can be configured with complex Boolean metacode filters allowing customizable news profiles. Stories are accessed by double clicking the appropriate headline.

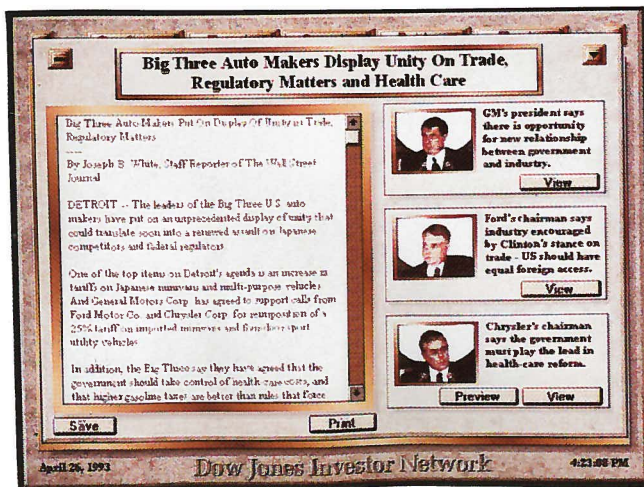


Plate 3

A Multimedia story is organized into a collection of hierarchical viewers. This top level story tableau contains a simple Wall Street Journal text viewer, and three Icon Select Panels which allow access to separate sub-story media viewers. There is also the capability to provide a preview clip (media bite) of selected material.

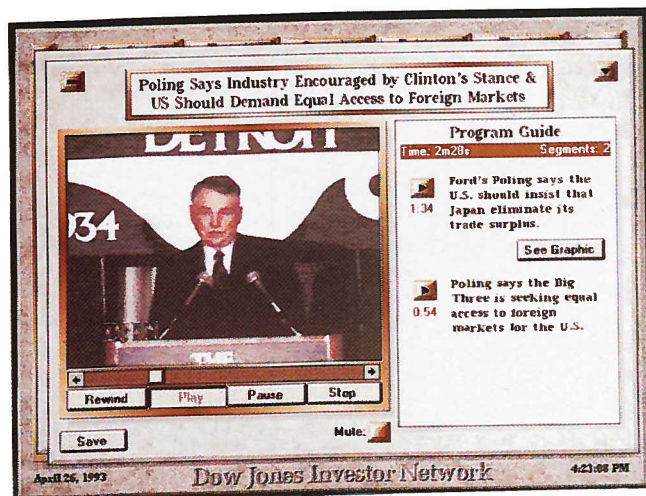


Plate 4

A typical low level media viewer allows several techniques to quickly access media information: standard VCR-like motion controls, and random playback re-positioner (horizontal slider). The Program Guide allows non-linear access directly to pre-indexed scenes, as well as hyperlink references to other media viewers or stories.