Open Blueprint

Multimedia Resource Manager

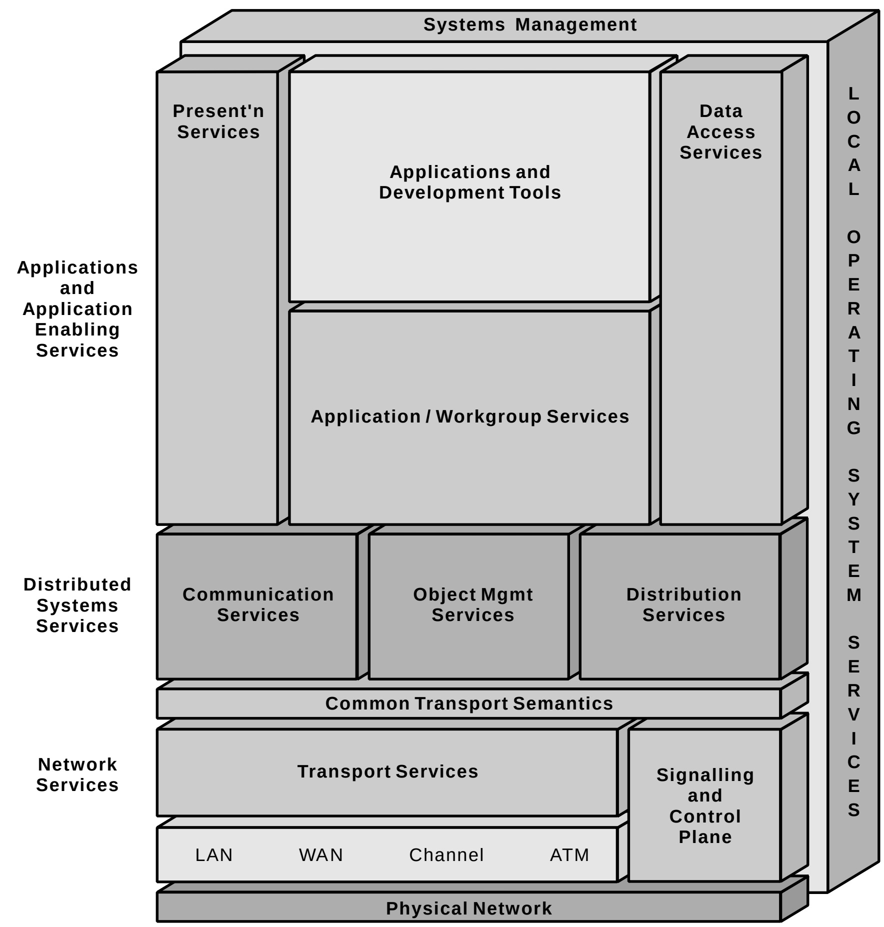

GC23-3921-01

Open Blueprint

Multimedia Resource Manager

GC23-3921-01

About This Paper

Open, distributed computing of all forms, including client/server and network computing, is the model that is driving the rapid evolution of information technology today. The Open Blueprint structure is IBM's industry-leading architectural framework for distributed computing in a multivendor, heterogeneous environment. This paper describes the Multimedia resource manager component of the Open Blueprint and its relationships with other Open Blueprint components.

The Open Blueprint structure continues to accommodate advances in technology and incorporate emerging standards and protocols as information technology needs and capabilities evolve. For example, the structure now incorporates digital library, object-oriented and mobile technologies, and support for internet-enabled applications. Thus, this document is a snapshot at a particular point in time. The Open Blueprint structure will continue to evolve as new technologies emerge.

This paper is one in a series of papers available in the *Open Blueprint Technical Reference Library* collection, SBOF-8702 (hardcopy) or SK2T-2478 (CD-ROM). The intent of this technical library is to provide detailed information about each Open Blueprint component. The authors of these papers are the developers and designers directly responsible for the components, so you might observe differences in style, scope, and format between this paper and others.

Readers who are less familiar with a particular component can refer to the referenced materials to gain basic background knowledge not included in the papers. For a general technical overview of the Open Blueprint, see the *Open Blueprint Technical Overview*, GC23-3808.

Who Should Read This Paper

This paper is intended for audiences requiring technical detail about the Multimedia Resource Manager in the Open Blueprint. These include:

- Customers who are planning technology or architecture investments
- Software vendors who are developing products to interoperate with other products that support the Open Blueprint
- Consultants and service providers who offer integration services to customers

Contents

Figures

Summary of Changes

This revision describes:

- The consolidation of Open Blueprint multimedia infrastructure support into the Multimedia resource manager
- The support provided by the Multimedia resource manager to exploit Internet connectivity
- Media streaming support

Multimedia Resource Manager

Distributed multimedia systems involve a number of services to handle various kinds of data, and applications that coordinate the presentation of this data to provide the user with a multimedia experience. Strictly defined, the term *multimedia* implies two or more forms of digital media, at least one of which is a time-based digital medium such as audio, animation, or video. The data that is not time-based, however, is usually handled by other resource managers. Multimedia applications interact with several resource managers to combine the different data types to present a multimedia experience to the user. Because time-based digital data such as audio or video is particularly important for multimedia applications, it is often called *multimedia data*. This description of the Multimedia resource manager describes a strategic architecture and the interfaces for time-based digital data services.

Time-based digital data, or multimedia data, when stored, can occupy considerable storage space. When recorded or played back, it must appear as a continuous stream of data packets, or *data stream*, which must be delivered from source to destination within stringent performance constraints. For example, the data must be delivered from the disk drive on a file server across a network and up to the user interface. Insuring that these constraints are met in a distributed system environment is known as guaranteeing the *quality of service* (QoS).

The presentation services for multimedia applications are usually implemented on personal workstations. Each of the major workstation operating systems has evolved its own architecture for processing and presenting multimedia data. The collection of functions implementing this architecture is usually referred to as multimedia system services. The Multimedia resource manager described in this document interacts with and supplies multimedia data to these local multimedia system services.

The Multimedia resource manager consists of core services which form the foundation upon which enhanced multimedia services are built. The enhanced multimedia services extend the core by customizing it for particular software infrastructures or application domains. The client portion of the Multimedia resource manager is called the *viewer*. It provides the interfaces to using applications.

The next section describes the Multimedia resource manager core services. The following two sections present the Internet video services and the media streaming services, two enhancements of the core. A section on the workstation application environments briefly sketches the multimedia system services used by the application. The final section gives a brief introduction to related resource managers.

Multimedia Resource Manager Core Services

The Multimedia resource manager records, manages, retrieves and plays time-based digital data such as video and audio streams. It views this data as continuous digital data streams to be delivered at a given data rate. The Multimedia resource manager's services are independent of specific data stream attributes. In particular, the architecture is independent of the compression method employed, the stream encapsulation method (such as MPEG-2 transport stream), and the network protocol used to transmit the stream to the application. Where these stream attributes need to be considered for the stream management, an extension structure allows the insertion of stream related modules. This architecture isolates functions related to specific data streams from the generic functions, and permits the accommodation of new stream types as needed.

The Multimedia resource manager manages the data that forms the real-time streams, and the *metadata* it needs for their delivery. Examples for such metadata are the compression method employed, the streaming data rate, and the size of the files. A multimedia *asset* consists of the data stream and the associated metadata. Information about the content of the data streams, such as actors in a movie, copyright information on a news clip, or billing information, is managed by an associated application. Other services required to build a solution, such as database services, digital library services, or transaction services, are supplied by their respective resource managers. The Multimedia resource manager is designed to support many application types, from digital video broadcast applications over news-on-demand systems to video serving on the Internet.

Architecture

The Multimedia resource manager *core services* contain the building blocks from which one can build a wide range of multimedia services configurations and application systems. The Multimedia resource manager architecture is defined to be scalable from small single system servers to very large complexes serving thousands of streams. Figure 1 below illustrates an overview of the core services. The major components of these services are the *control services* and the *data pumps*. A data pump consists of the real-time file system, the data exporter, and network access routines. These components are present in any implementation of the Multimedia resource manager. Archive management provides off-line storage that is accessible through the Multimedia resource manager, but without real-time delivery capabilities.

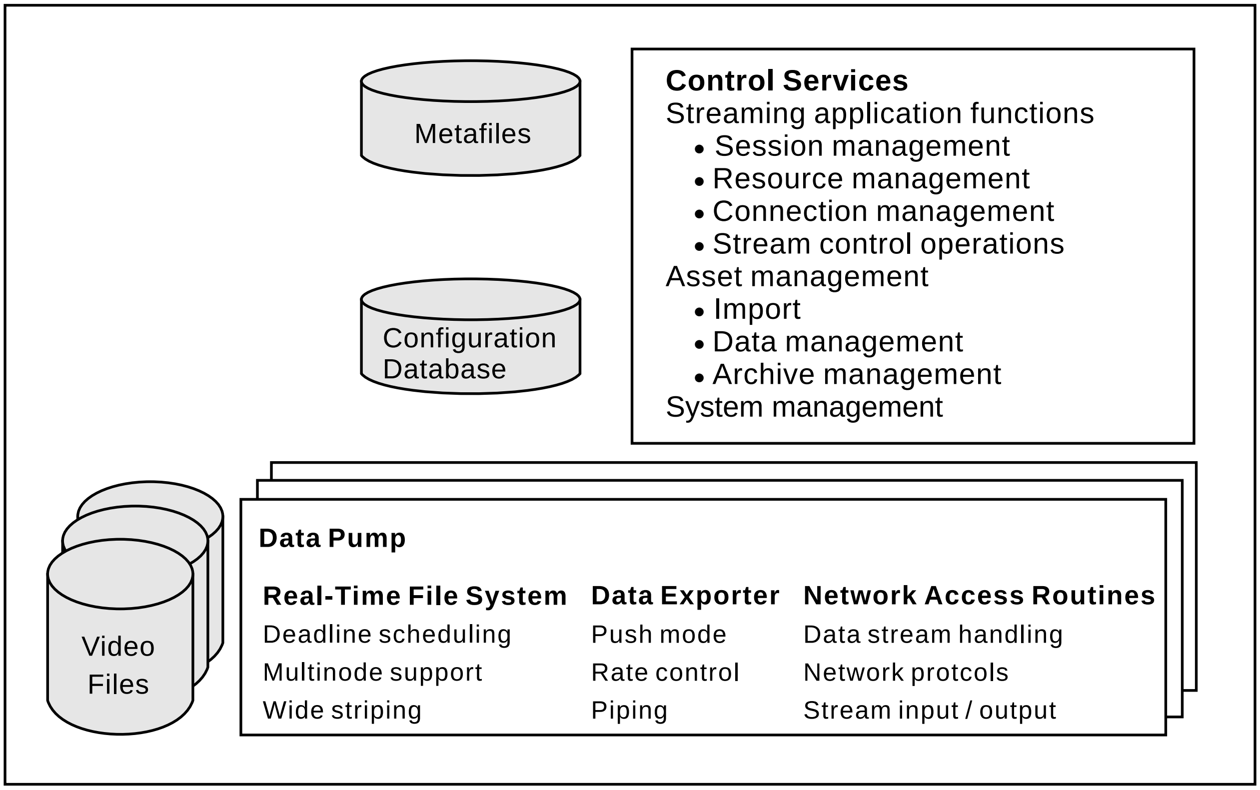

Figure 1. Multimedia Resource Manager Core Services

To guarantee quality of service, the Multimedia resource manager must account for resource usage on the system on which it is implemented. For instance, the Multimedia resource manager must know how much processing power and I/O bandwidth are currently used, how much it is capable of, and how much an additional data stream would consume. Based on this bandwidth management function, the Multimedia resource manager performs *admission control* when each new request is received to determine whether the Multimedia resource manager has sufficient bandwidth resources to support this stream. If there are sufficient resources, admission control reserves the resources and admits the request. Otherwise the request is denied. The knowledge of dynamic resource consumption in the Multimedia resource manager requires that it operates in a controlled environment. Also, serving time-based data is resource intensive. Therefore, the Multimedia resource manager core services are usually implemented as a server in a client/server environment. The other components of a multimedia application can reside on servers or on client systems.

Control Services

The *control services* act as a single control point for importing, managing, and retrieving and playing multimedia assets. To meet this objective the control services provide several classes of functions.

- The streaming application functions support retrieval and playing of data streams. They provide an application with the ability to establish a session with the Multimedia resource manager. Resource management functions allocate real-time resources, such as the processing power and the I/O bandwidth of the Multimedia resource manager, and perform admission control. The connection management functions establish and manage the connections from the data pumps to the viewers. The stream control operations allow the applications to control real-time data streams.
- The asset management functions support content management applications. They perform data import and data management, including data placement across multiple data pumps. Metafiles contain the assets' metadata such as the data rate at which a stream needs to be played, the type of

compression used, and the data pumps on which the data files are located. Resource management and connection management use the data in the metafiles for their operations.

The asset management functions also include archive management which interfaces to the Storage Management resource manager. Archive management provides off-line storage for a large amount of video data. The archive can be configured on a separate server. In this case, only the asset management functions of the control services are present. No real-time stream delivery functions are provided. The data from an archive is loaded onto a data pump using the stage commands between their respective control services.

- The system management functions are used for administering and monitoring the Multimedia resource manager. Administration functions define the hardware and software configuration of the Multimedia resource manager and its connection to the network. Examples of these configuration data are the number of data pumps, their storage capacity and bandwidth capacity, and the number and addresses of network interfaces supporting a particular network protocol. Administration functions are accessed through HTML panels in a Web browser. Dynamic monitoring is performed from an external workstation through a management application using the SNMP protocol.

The control services provide a single, uniform image of the Multimedia resource manager to its applications, hiding from them the configuration and API details of the data pumps. Where configuration details cannot be hidden completely, they are exposed through logical concepts that can be mapped to physical components in a way that is transparent to the application. In particular, not all data pumps in a resource manager can contain the same assets. This requires that the content management functions know about grouping of assets onto data pumps. However, rather than exposing the configuration of the data pumps to the content management application, the Multimedia resource manager externalizes the logical concept of an *asset group*. The resource manager can map asset groups to data pumps transparent to its applications. Similarly, if not all data pumps have network connections to all potential clients, the resource manager does not directly expose the network configuration. Instead, it externalizes the concept of a *port group* to the application, allowing for transparent mapping of port groups onto the actual network configuration.

The control services are also a focal point for resource optimizations within the server complex that implements the Multimedia resource manager. Examples of optimizations are placement of data files onto data pumps to balance the workload, and selection of data pumps from which to play a stream when multiple data pumps hold a replica of that stream.

Data Pump

A data pump stores and delivers audio and video streams. It has a real-time file system that is optimized towards delivering time-based digital data streams. The files for such data are typically large, are read sequentially, and have stringent timing requirements for delivery. Due to their size, most of the data streams are compressed. The compression algorithms employed usually produce streams with a fixed data rate, regardless of the characteristics of the stream contents.

The data exporter acts as the pump, reading data streams from the real-time file system and pushing them at the required rate into the network. This push mode requires interactions between the viewing application and the Multimedia resource manager only when the application requests a change in a stream. When a client requests a stream to be played, the server maintains a smooth data flow into the network. The data exporter receives its control commands from the control services.

The data exporter's piping function takes as input a video stream transmitted by another data pump and sends this stream to an analog decoder. This feature can be used connect media streaming systems. For instance, a Multimedia resource manager might not hold any videos on disk storage, but is installed on a system that has video decoder cards to generate analog video streams. In this case, the Multimedia resource manager receives its data streams from a remote Multimedia resource manager over a network such as ATM. This configuration allows central storage of the data streams and use of ATM for wide-area networking, exploiting the attendant economies. At the same time, it can exploit a cost-effective local analog distribution network.

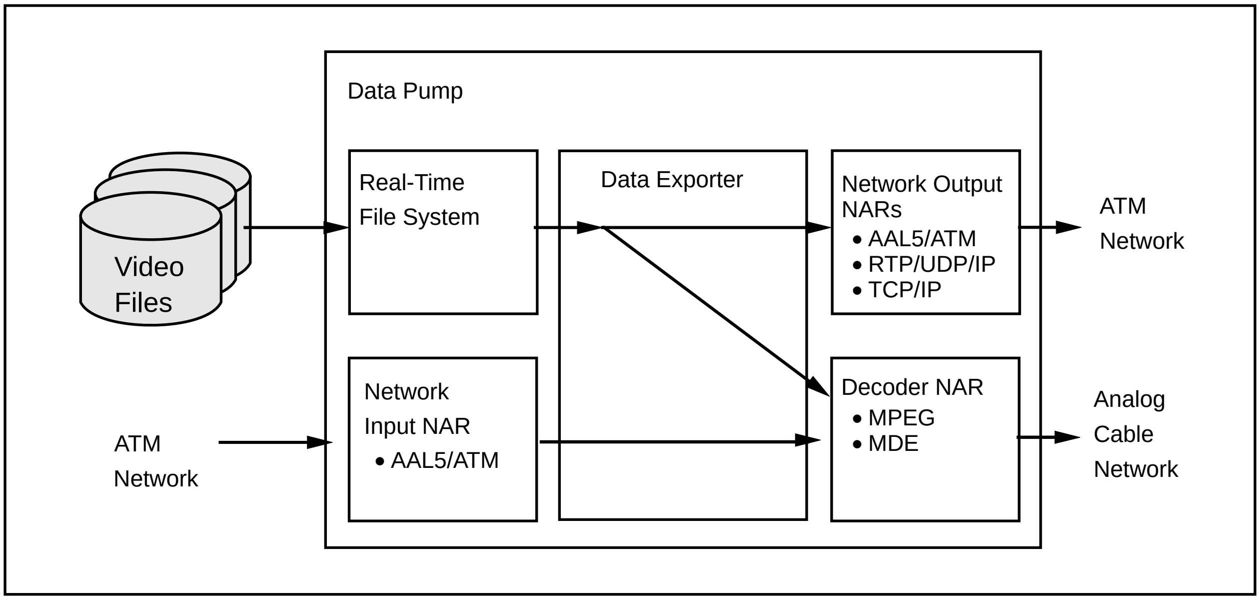

Figure 2. Data Streaming Configurations

The network access routines (NARs) are stream-related modules that perform operations specific to particular types of data streams or to particular network protocols. Some examples of network protocol suites supported by particular NARs are:

- MPEG-2 Transport Stream over AAL5 and ATM
- MPEG-2 Transport Stream to built-in decoder card with analog output
- Any video or audio stream point-to-point over RTP/UDP/IP
- Any video or audio stream multicast over RTP/UDP/IP

For non-real-time input, the data pump supports the TCP/IP LAN/WAN protocol using the File Transfer Protocol (FTP) application protocol. For real-time input to piping operations, the data pump supports MPEG-2 transport streams over AAL5 and ATM.

Data Stream Dependent Extensions

Some data stream types have special structural attributes that the *Network Access Routines (NARs)* need to consider when transmitting the streams. Examples of such attributes are the configuration tables in MPEG-2 transport streams. To limit knowledge of these attributes in the Multimedia resource manager, the attributes are stored transparently in the control services' metafiles, and delivered to the NARs without interpretation by either the control services or the data exporter. NARs can be added to the data pump as new stream types become available, allowing for the adaptation of the resource manager architecture to new stream encapsulation methods.

Scalability and Availability

The Multimedia resource manager can be implemented with a wide range of storage and delivery capacities. At the entry of this range, all components, including the applications, can reside on a single system. Next, the application and the Multimedia resource manager can reside on different systems. A single instance of the control services can support many data pumps that have access to a common disk storage pool or separate disk storage pools. The only limit is the ability of the control services to process the control transactions issued by the application. If a single system cannot support a given rate of control transactions, the control services can be implemented on multiple systems, while maintaining a single, un-partitioned image of the Multimedia resource manager.

The fact that a single image of the Multimedia resource manager can contain multiple instances of each of its components also aids in its recovery functions. For instance, upon the failure of one instance of the control services, another instance can take over without interrupting the data streams being transmitted by the data pumps.

Programming Interfaces

Application Interfaces

The control services provide all explicit application programming interfaces to the Multimedia resource manager. These interfaces can be divided into the streaming application interfaces and the asset management interfaces.

The streaming application interfaces support session management functions, connection management functions, and stream control operations. The session management functions control the relationship between the application and the Multimedia resource manager. A session provides an operating context, end points for communication, and a grouping mechanism for related data streams. The connection management functions control the real-time data connections over which the data pump receives and sends the data streams. The stream control operations allow the application to control the data streams themselves, with commands such as PLAY, PAUSE, and JUMP.

The asset management interfaces support the import and management of multimedia assets. Importing an asset includes selecting the data pump(s) on which to place the data file, copying the data file to the data pump, and setting up the meta data on the control services. Other asset management functions can change asset attributes, move assets to and from the archive, and delete assets.

The streaming applications and the content management applications communicate with the control services using DCE RPCs. The Multimedia resource manager provides a library for installation on the application systems. Using the library calls, these applications can easily implement the calls to the Multimedia resource manager and are isolated from the details of implementing DCE RPCs through the Interface Definition Language (IDL).

System Management Interfaces

The *system management interfaces* support the management of the Multimedia resource Manager's hardware and software configuration and their connection to the network.

Relationship to Other Solution Components

The Multimedia resource manager does not by itself represent a solution. It must be controlled by an application, and embedded into a broader infrastructure that provides content preparation and management, and a network infrastructure. Figure 3 on page 11 shows how the Multimedia resource manager fits into an end-to-end solution and how it connects to the other components of that solution.

Two types of applications use the Multimedia resource manager:

- **Content Management Applications**

  Content management applications use the asset management functions to manage the real-time data. At the same time, they can store other data related to the audio and video streams, such as authors, actors, and billing information, on a content management data base. An example of a content management application is the Digital Library resource manager.

- **Content Delivery Applications**

  Content delivery applications (applications that deliver video and audio streams) connect the Multimedia resource manager to the remainder of the delivery infrastructure. These applications can reside on the client or server system. Examples are Internet video applications, a web browser, a bridge to a video studio automation system, or a scheduler application to schedule broadcasts and Near-Video-on-Demand (NVOD) streams.

The viewer is the component where the data streams are rendered for viewing by people. The viewer can reside on a workstation or on a television set-top box. The viewer receives the data stream over a high-bandwidth network. It communicates with an application to exercise this control rather than directly with the control services. The application can reside on the system where the viewer is located, as in the case of a workstation, or on a special application system, for example, when the viewer is a television set-top box. The viewer often has very limited control over the video and audio streams. In applications such as broadcast services or with NVOD, the viewer can only select a particular channel. All other kinds of controls are exercised by a scheduling application residing at the broadcast studio.

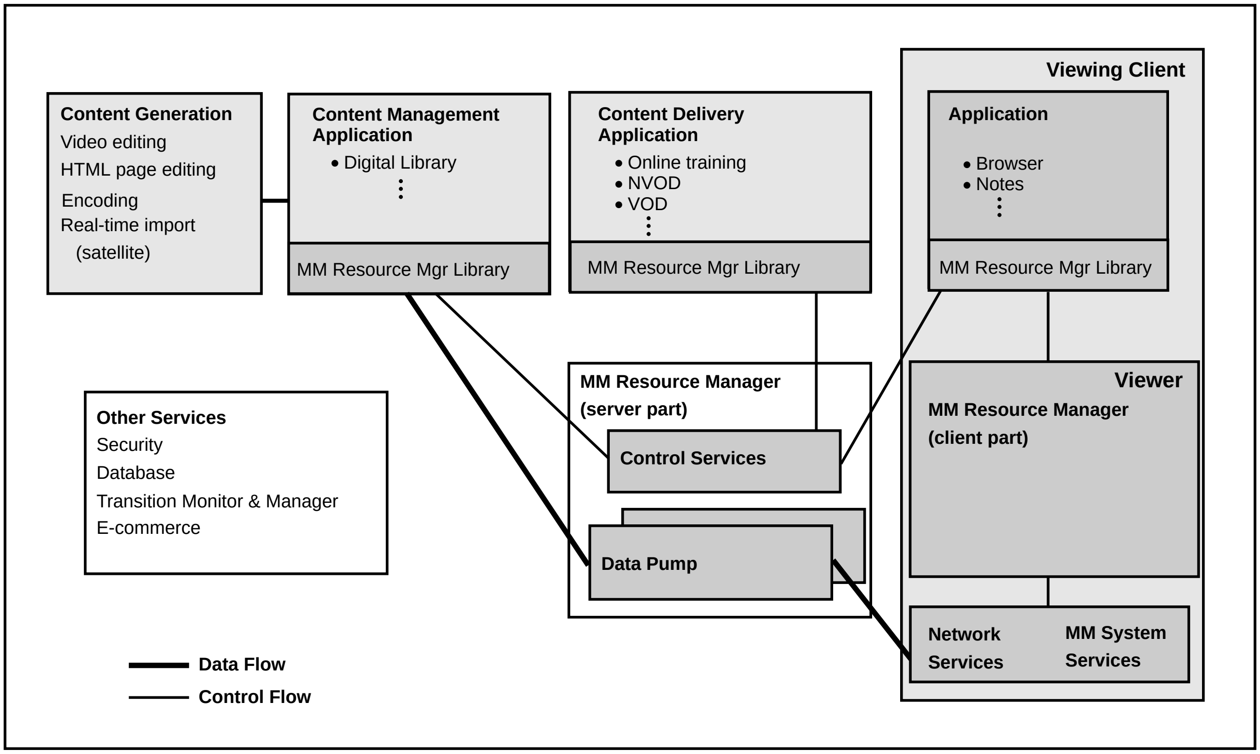

Figure 3. Multimedia Resource Manager and Other Solution Components

Internet Video Services

The Internet video services augment the store-and-forward information delivery prevailing on the World Wide Web with data streaming. The Internet video services are built as extensions of the Multimedia resource manager core services. They deliver streaming video and audio to workstations connected to the Internet or to an intranet. With data streaming, the server pushes data at a continuous rate, and the client can start consuming the data as soon as it gets a required minimum amount. In particular, the client does not have to wait until it has received all the data. Data streaming is particularly well suited to real-time data such as audio and video. The Internet video services are integrated into a World Wide Web infrastructure, providing their function in a well defined client/server framework.

The Internet video services use Internet protocols, such as HTTP for the client/server communication, HTML for page formatting, and TCP and UDP over IP as network protocols. The Internet video services can be deployed both on the public Internet, and on private intranets in a somewhat more controlled environment. The public Internet provides ubiquitous (World Wide!) access, but usually also implies low available bandwidth, such as 28.8 kbps achievable over a telephone connection. Corporate intranets often use LAN technology in a controlled and well engineered environment, and can supply much higher data bandwidth.

Users can use the Internet video services to augment existing Web applications with audio and video data. Furthermore, entirely new applications can be created that include audio and video while using the ease of development and the ubiquitous access of the World Wide Web.

Internet Video Services Structure

The Internet video services structure shows a clear distinction of client and server components. On the server, several specialized components are interconnected in modular fashion that provides standard interfaces where possible, and maximum flexibility to build highly scalable server configurations. The server consists of:

- The presentation formatter
- The content management application and database
- Video streaming applications
- The core services (control services and data pump)

The HTTP resource manager is used for communications with the clients.

On the client, the viewer is part of the Multimedia resource manager. It is invoked by the Web browser, or some other application, and it uses the network services and the multimedia system services of the client operating system.

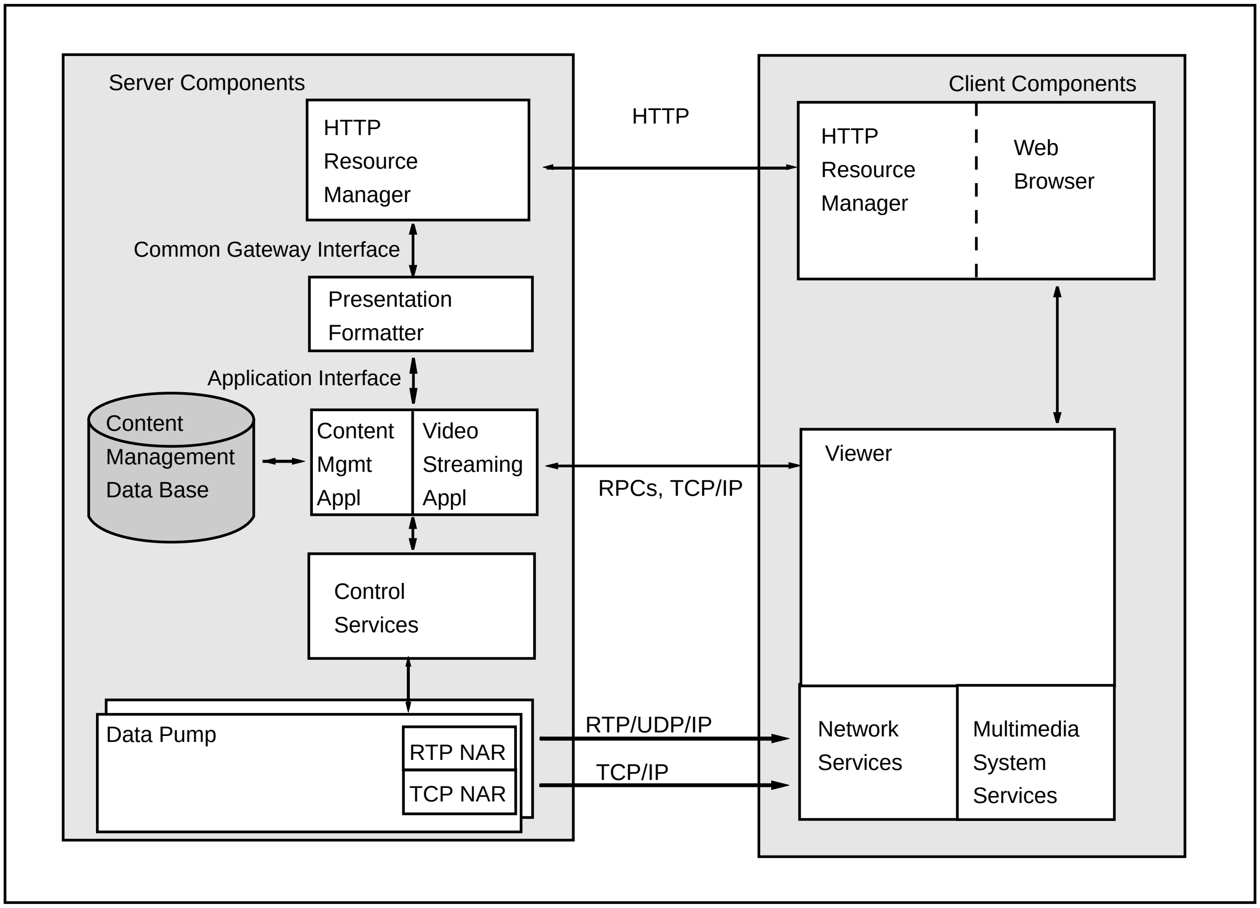

Figure 4. Internet Video Services Structure

Internet Video Services Components

Presentation Formatter

The *presentation formatter* consists of a set of HTTP Common Gateway Interface-binary (CGI-bin) programs that control the navigation and the selection of multimedia content. These programs are invoked through HTML forms the end user completes in the browser. The browser uses the HTTP resource manager to send the forms to the presentation formatter. The CGI programs are stateless programs that read information, process the information, pass HTML data back to the HTTP resource manager, and then exit.

The presentation formatter is the software platform on which a user can develop specialized streaming applications including video and audio. A client library supplied by the Multimedia resource manager makes the video application functions available in an easy-to-use form. Associated with the Multimedia resource manager are CGI programs implementing a video selection application, and a video-on-demand application. Users can tailor the presentation formatter to their applications by supplying their own CGI programs.

The presentation formatter also supplies the CGI programs for content management of the Internet video services. The user interface for content management is a Web browser. It invokes the content management functions by using the HTTP resource manager, which in turn calls CGI programs that use the application services' client library to perform content loading and content management.

Content Management Application

The content management application is responsible for providing the following functions:

- Content loading and management
- Navigation support using the content management data base
- Logging user transactions

The content management application provides a database for the storage and retrieval of information regarding the multimedia content stored on the Multimedia resource manager. This information is defined by the user and is used to present helpful information for end-users to make a selection of which content they would like to view. For example, if the content is comprised of various training videos, the information stored in the database might be subject, training prerequisites, duration, author, and so on. The content management application services provide the ability to query the database based on certain search criteria and the search results are then passed back to the caller (in this case the presentation formatter). The Digital Library resource manager provides content management application function (see “Internet Video Services Digital Library Integration” on page 19).

Content Management Database

The content management database holds information on the material stored in the data pump. Examples of such information are the director and actors of a movie, the copyright owner of a video, payment conditions, the date the video was created, and so on. In the context of the Digital Library resource manager, the content management database is the catalog portion of the library server.

Control Services

With the Internet video services, use of the control services interfaces is reserved for use by the content management and video streaming applications. Only the content management application's APIs are externalized for use by users to write applications, implementing them as HTTP CGI functions. The control services perform their function as described in “Control Services” on page 6.

Data Pump

The data pump stores and delivers the audio and video streams over the Internet to the clients. The general architecture of the data pump is described in the previous section.

The preferred data connection to the clients uses the Real-Time Protocol (RTP) that flows over UDP/IP. The TCP/IP protocol is an alternative to the RTP/UDP/IP stack. The data pump supports point-to-point (unicast) video transmission, and, over the RTP/UDP/IP protocol, multicast transmission. For network resource reservation, the data pump supports the Resource Reservation Protocol (RSVP). These protocols are implemented in specialized NARs.

Viewer and Video Streaming Application

The viewer controls the data flow on the client and the operation of the multimedia system services. The viewer provides a graphical user interface for controlling the video streaming applications. The viewer connects to the video streaming application using a network connection. Over this connection, the viewer transmits VCR-like commands such as PLAY, PAUSE and SEEK to control the data pump. The viewer also sets up the network services on the client so that it can receive the video stream over the data connection and pass it to the multimedia system services for decoding and rendering. The viewer, portions of the Open Blueprint Network Services, and the multimedia system services are integral to the Internet video services.

The video streaming application interacts with the control services to manage video streams. Video streaming applications have hooks through which a user can invoke billing or accounting functions.

Related Solution Components

HTTP Resource Manager

As the front end to the system, the HTTP resource manager is responsible for receiving and processing the HTTP commands. Navigation through the video application and the selection of videos to play occur through interactions between the HTTP services and the Web browser on the client. The HTTP services invoke CGI-bin scripts to determine the available videos by querying the presentation formatter, which in turn obtains this information from the content management application. The presentation formatter formats this information into HTML pages and metafiles, which the HTTP resource manager passes to the browser.

The HTTP services also log information about the incoming requests such as host name, date and time, action requested, success or failure, and number of bytes transmitted. The HTTP services uses the HTTP *forms* mechanism to obtain input from the browser, and it uses the Common Gateway Interface (CGI) to interact with the presentation formatter. Examples of software suitable for this function are the IBM Internet Connection server for AIX, and the Netscape server for AIX.

Web Browser Resource Manager

A standard browser provides the user interface on the client for navigation and selection of videos. When the user has selected a video, either through a simple click on a hyperlink or through a series of interactions with the server, the HTTP resource manager sends a metafile to the browser. Based on the MIME type of that metafile, the browser invokes the viewer, which initiates the transmission of the selected video. The client can use any standard Web browser software for this component. Examples of suitable browsers are the Netscape Navigator and the Microsoft Internet Explorer.

As a simpler alternative, which does not require a presentation formatter or content management application, the Multimedia resource manager supports serving video clips identified by Uniform Resource Locators (URLs) as used with the World Wide Web. This support is most useful when the viewer is a streaming plug-in to the browser. A streaming plug-in will start displaying a data stream as soon as it has received an initial portion, rather than waiting for the entire file identified by the URL.

The control services recognize a subset of HTTP. In particular, they interpret a GET function as request to view a video. They select a data pump, and send the video request to it. Then they use the redirection feature of HTTP to cause the browser to connect to the selected data pump. When the browser connects to the data pump, it streams the selected video stream over the TCP/IP connection. The browser feeds the stream directly to the streaming plug-in which decodes and renders the video. Figure 5 provides an overview of the Web browser support of the Multimedia resource manager.

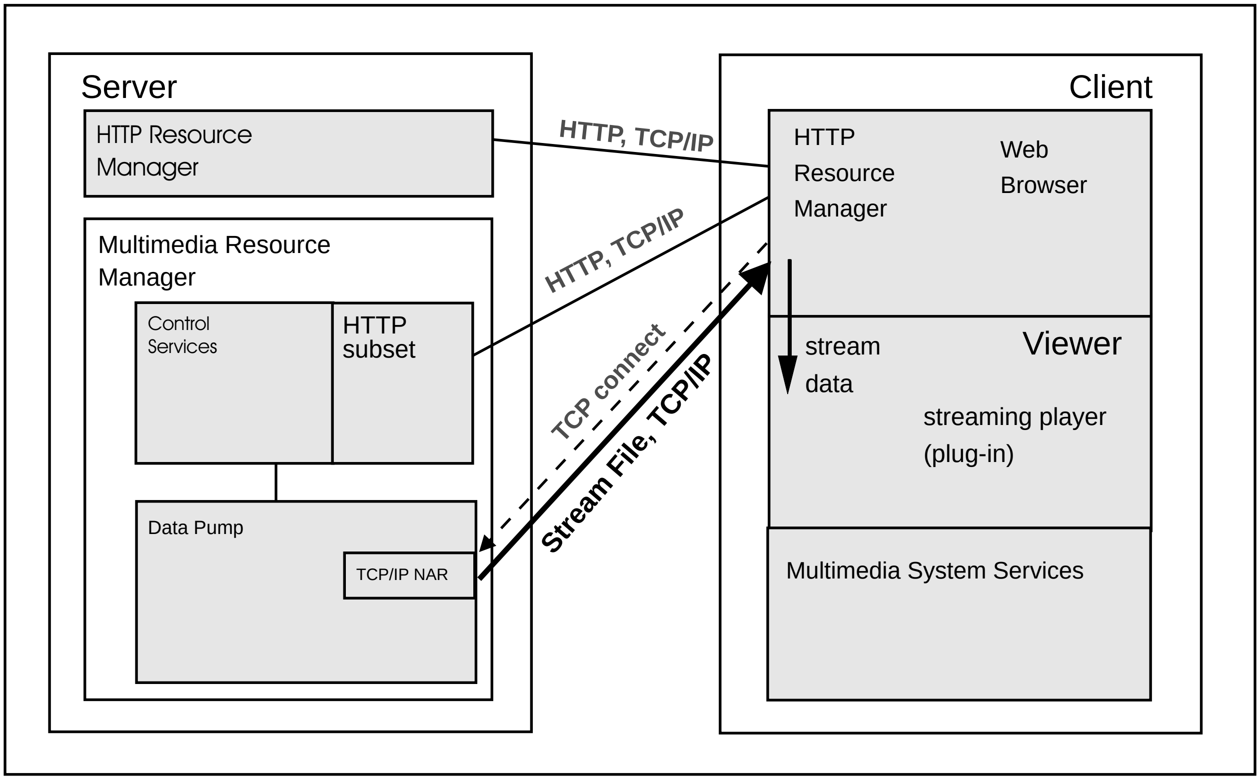

Figure 5. Overview of the Web Browser Support of the Multimedia Resource Manager

Network Services

The network services manage the reception of the video stream from the network. They use UDP and IP as the transport protocols. The RTP protocol on top of UDP/IP employs time stamps and sequence numbers to repair transmission defects if possible before the network services pass the video stream to the multimedia system services. Where available, the network services use the Resource Reservation Protocol (RSVP) to communicate network bandwidth reservation requests to the IP network.

Multimedia System Services

The multimedia system services decode and render the compressed video and audio streams. The Multimedia system services support several compression algorithms. MPEG and AVI decompression are supported for networks with high bandwidth capabilities, such as intranets. The client software can use either a software decoder, such as the decoders shipped with Microsoft Active Movie, or an MPEG hardware decoder when one is available. An H.263 video decoder and an G.723 audio decoder are included. These decoders are suitable for low bit rate transmissions over telephone lines.

Internet Video Services Interface

The following sections describe the interfaces a user or an integrator can use to integrate the Internet video services into a wider application framework.

Application Interface

The content management and video streaming applications and their associated library provide the programming interface for the presentation formatter CGI programs.

- Session establishment between the presentation formatter and the content management application
- Video selection
- Creation and transmission of HTML metafiles
- Content management database functions such as query, add, and delete
- Asynchronous callbacks for user supplied actions like billing or accounting
- Content loading

Internet Video Services Digital Library Integration

The Digital Library resource manager can fulfill the function of a content management application. The Digital Library resource manager provides a full range of content creation and capture; storage and management; distribution; search and access; and rights management support for collections of digitized information such as text, the bit maps of images, or digitized video. The digital library server supports a catalog and the collected content that is managed by the Digital Library resource manager.

In an Internet video services context, the Digital Library resource manager catalog is the content management database described above. Video files that are part of the digital library collection can be stored in the Internet video services data pump real-time file system; requests to play them are satisfied by requesting the Multimedia resource manager to do so. Video files that are stored elsewhere in the Digital Library resource manager collection must be copied to the Internet video services data pump real-time file system before they can be streamed.

For a complete description of the Digital Library resource manager, refer to the *Open Blueprint Digital Library Resource Manager* component description paper.

Media Streaming Services

The television industry is in the process of changing from analog technology to digital technology for storage and delivery of video. This changeover occurs gradually as digital video technology becomes available and cost effective. At the same time, digital video technology opens new opportunities for video distribution. Examples of applications are news distribution and replacement of video tape recorders in television studios.

The media streaming services support a wide range of applications. These applications generally use television sets to present the streams to users, rather than personal computers and workstations. The applications range from automating video broadcast studios and cable television head-ends to digital video editing, custom news distribution systems and Near-Video-on-Demand (NVOD) systems. The data streams can be distributed digitally over ATM networks, or in analog or digital form over cable television networks. A hybrid of the two distribution modes is digital television distribution over a cable television infrastructure. As the technology matures and the end user infrastructure is upgraded, these applications could evolve into Video-On-Demand (VOD) applications.

Media Streaming Structure

The media streaming services record, manage, and deliver digital video streams to the network. Additionally, the media streaming services can also decode compressed video files and deliver them in analog form over cable networks.

The media streaming structure supports video servers with great flexibility in their input and output options, with quality of service (QoS) assured throughout, and scalability from small systems to very large systems. Figure 6 on page 23 shows an overview of the media streaming structure. The media streaming services fully utilize all components of the Multimedia resource manager core services, including the archive function and the piping function.

The major components of the media streaming structure are:

- The control services
- The data pump
- The distributor

The control services and one or more data pumps are central to a media streaming system. The other components, such as the distributor or a Storage Management resource manager for archiving, are added as required. The initial support is for smaller systems. Larger system complexes can be built by adding more data pumps to form a server cluster. Even larger systems, based on implementing the Multimedia resource manager on a parallel supercomputer, will evolve as the demand develops.

Control Services

In the media streaming structure, the control services' interfaces are directly externalized for users and integrators to write their applications. The control services perform their functions as described in the section on the core services of the Multimedia resource manager.

Data Pump

In the media streaming structure, the data pumps contains its regular components, the real-time file system and the data exporter. In addition, data pumps can use special hardware to support decoding of compressed video streams, or multiplexing of digital video streams for cable television networks.

Decoder hardware can decode a wide variety of compressed data streams as required by the television industry. In many cases, the data is compressed with the MPEG-2 compression algorithm and are encapsulated in an MPEG-2 transport stream that allows for multiplexing and synchronizing multiple time-based data streams. The decoder hardware selects a video and an audio stream from the multiplex, decodes them, and sends them in analog form over a cable television network. Depending on the application, different decoders might be needed to handle different compression profiles from low bandwidth to high quality video streams.

Multiplexing hardware can combine multiple MPEG data streams into a single MPEG-2 transport stream and transmit that stream in digital form to a cable head-end or to a set-top box for decoding. Depending on the application, multiplexors must handle a range of data stream bandwidths and insertion of real-time controls into the transport stream.

Data stream-dependent extensions are supported by their respective specialized NARs. These extensions include:

- MPEG-2 transport stream over ATM
- Decoding in the data pump
- Analog transport
- Multiplexing MPEG-2 transport streams over digital cable networks

When needed, the asset management functions in the control services support storing of metainformation specifically for use by the NARs. When the control services request a data pump to start a data stream with metainformation for the NAR, it passes this information through the data exporter to the NAR. This metainformation is special to the NAR, and neither the control services nor the data exporter read or interpret it.

Media Streaming Distributor

In addition to the piping function of a data pump, the media streaming services have a special purpose distributor component. This media streaming distributor is a hardware component that accepts the MPEG-2 transport stream over ATM as input. Using its resident decoder cards, the distributor converts the MPEG-2 transport stream into analog video streams.

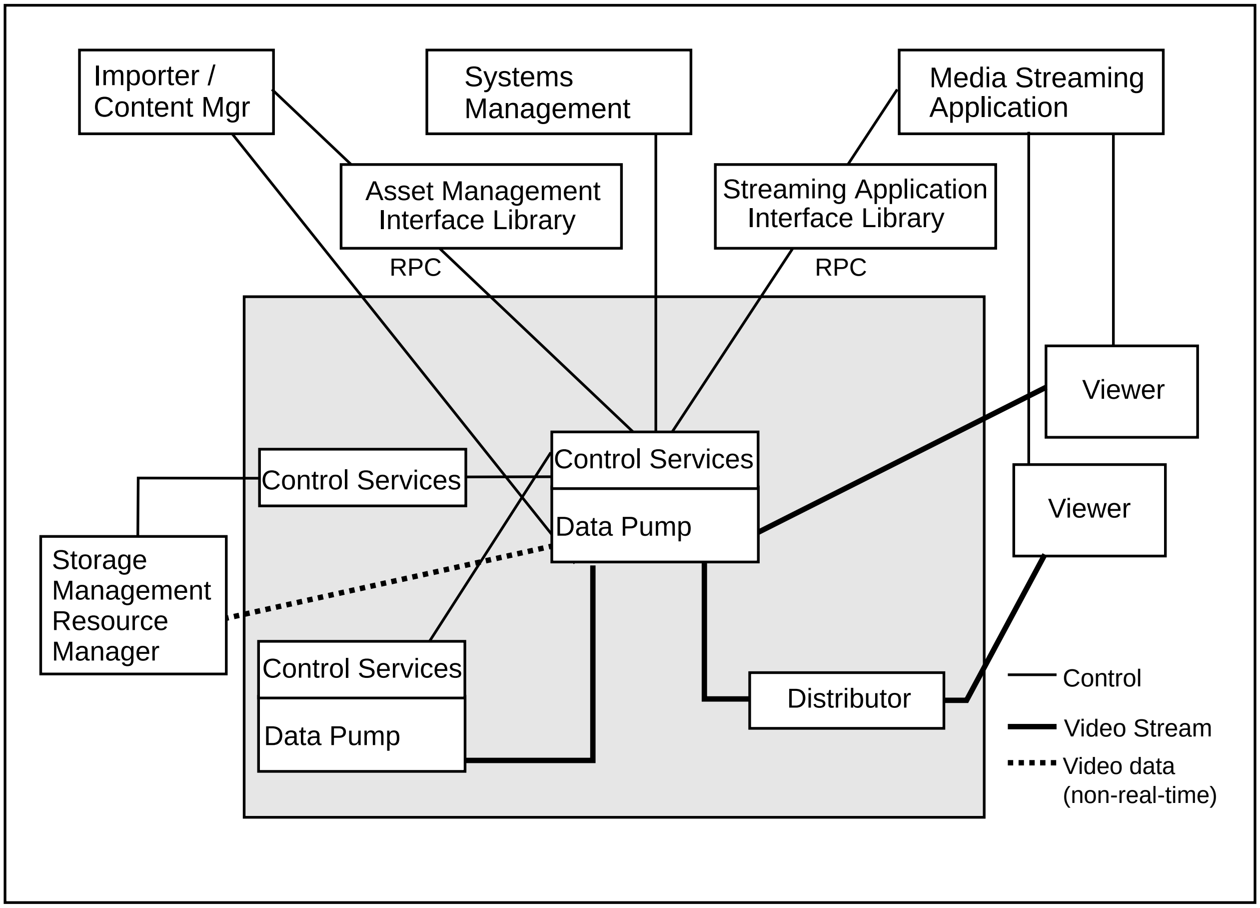

Figure 6. Media Streaming System Structure

Media Streaming Services Interfaces

Media streaming applications directly use the control services programming interfaces. These interfaces are described in “Programming Interfaces” on page 9.

Related Solution Components

Application

A media streaming application is needed to connect the media streaming services to the remainder of the video infrastructure. The media streaming application is a bridge to an existing studio automation system, or a scheduler application to schedule broadcasts and NVOD.

Media Streaming Viewer

The media streaming viewers are the components where the video stream is rendered for viewing by people. In media streaming environments, this is usually a television set, connected to the video streaming infrastructure through a set-top-box. The viewer usually has very limited control over the video streams. It communicates with the application to exercise this control rather than directly with the control services. In many applications, however, such as with broadcast services or with NVOD, the user can just chose whether to watch a particular channel. All other kinds of controls are exercised by a scheduling application residing at the broadcast studio.

Multimedia System Services

The model for Multimedia system services is typically found in a desktop operating system environment. Multimedia system services provide an interface to activate, control, and deactivate the resources needed to deliver a data stream from a source multimedia device to a destination multimedia device in the system that hosts the viewer.

A buffer pool in memory is used to exchange the data stream between source and destination. The interface is used by an application to control and synchronize data streams and to notify an application when an event has occurred. The application does not handle time-based media. It controls multimedia devices, data stream flow, and synchronization. The multimedia devices are assumed to be local, generic models of a class of devices, such as a VCR. Applications are shielded from the unique control characteristics of a vendor's multimedia devices. Multimedia system services is designed to be extensible, and this extensibility is used to add support for new classes of multimedia devices in successive operating system releases.

These services are very dependent on the operating systems on which they are implemented. Each of the major operating environments has evolved its own implementation. Multimedia resource manager interfaces can be invoked from viewers implemented on any of the major system platforms.

In Windows '95 and Windows NT systems, multimedia system support is provided as ActiveMovie support. A filter graph forms a pipeline for processing video and audio streams from the network adapters to the video and audio adapters. A filter graph manager coordinates the operation of the pipeline. The Internet video services support the ActiveMovie infrastructure and supply a network source filter for network attachment to the data pump, and H.263 and G.723 decoder filters for video and audio.

In the OS/2 system, Multimedia Presentation Manager/2 (MMPM/2) performs the functions of the multimedia system services. In AIX, Ultimedia Services provide the multimedia systems services functions. On the Macintosh, the QuickTime and QuickDraw support provides the multimedia system services functions.

Television set-top boxes typically have rudimentary operating systems that support only functions needed by the multimedia system services. An example for such an operating environment is the OS/9 system that is widely used in smart television set-top boxes.

Appendix A. Bibliography

DSM-CC: ISO/IEC JTC1/SC29/WG11 N100 rev1: Digital Storage Media Command and Control. Dallas, November 1995.

MPEG-2 Transport Stream: ISO/IEC/JTC1/SC29/WG11 N0801: Systems Recommendation H.222.0. November 1994.

MPEG-2 Video Encoding: ISO/IEC/JTC1/SC29/WG11 N602. Recommendation H.262. November 1993.

Introduction to MPEG: Didier Le Gall: MPEG: A Video Compression Standard for Multimedia Applications. Communications of the ACM, April 1991, Vol 34. No 4. pp. 46-58.

H.263 standard: ITU-T Draft H.263: Line Transmission of non-Telephone signals. May 1996.

ITU-T Recommendations, Coding of analogue signals by methods other than PCM: G.723 () Speech coders.

Open Blueprint Storage Management Resource Manager component description paper.

Open Blueprint Digital Library Resource Manager component description paper.

Multimedia Presentation Manager/2 Online Library, which is a component of the IBM Online Library, SK2T-2176.

Appendix B. Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Subject to IBM's valid intellectual property or other legally protectable rights, any functionally equivalent product, program, or service may be used instead of the IBM product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, is the user's responsibility.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
500 Columbus Avenue
Thornwood, NY 10594
USA

Trademarks

The following terms are trademarks of the IBM Corporation in the United States or other countries or both:

AIX
IBM
IBMLink
Multimedia Presentation Manager/2
Open Blueprint
OS/2
Ultimedia

The following terms are trademarks of other companies:

ActiveX	Microsoft Corporation
DCE	The Open Software Foundation
Netscape (logo)	Netscape Communications Corporation
Windows NT	Microsoft Corporation
QuickDraw	Apple Computer, Incorporated
QuickTime	Apple Computer, Incorporated

Microsoft, Windows, and the Windows 95 logo are trademarks or registered trademarks of Microsoft Corporation.

Appendix C. Communicating Your Comments to IBM

If you especially like or dislike anything about this paper, please use one of the methods listed below to send your comments to IBM. Whichever method you choose, make sure you send your name, address, and telephone number if you would like a reply. Feel free to comment on specific error or omissions, accuracy, organization, subject matter, or completeness of this paper.

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

- If you prefer to send comments by FAX, use this number:

  United States and Canada: 1-800-227-5088.

- If you prefer to send comments electronically, use one of these ID's:
  - Internet: **USIB2HPD@VNET.IBM.COM**
  - IBM Mail Exchange: **USIB2HPD at IBMMAIL**
  - IBMLink: **CIBMORCF at RALVM13**

Make sure to include the following in your note:

- Title of this paper
- Page number or topic to which your comment applies

Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.

GC23-3921-01