DAVIC
Digital Audio-Visual Council

DAVIC 1.5 Specification

Baseline Document 1

Revision 4.0

Applications for Home Storage and Internet Based Systems

Editor: - Henry Chadwick (hdchad@us.ibm.com)

Published by Digital Audio-Visual Council
Geneva, Switzerland

Table of Contents

1. INTRODUCTION 1

2. KEY CONCEPTS 2

2.1 SOURCE CONTENT IS LOCATED IN TIME AND SPACE 2
2.2 SELECTION 3
2.3 FULFILLMENT 4
2.4 MANAGEMENT AND USE OF ACQUIRED CONTENT 5

3. APPLICATION SCENARIOS 6

3.1 BROADCAST CAPTURE 7
3.1.1 Scope 7
3.1.2 Scenarios 8
3.1.2.1 User Initiated Capture 8
3.1.2.2 Agent Initiated Capture 8
3.1.3 Constraints and Comments 9
3.2 FILE TRANSFER FROM REMOTE STORAGE DEVICES 10
3.2.1 Scope 10
3.2.2 Scenarios 10
3.3 REMOTE STREAM ACCESS 10
3.3.1 Scope 10
3.3.2 Scenarios 10
3.3.2.1 Home Country from Hotel 10
3.3.2.2 Foreign Broadcast from Home 11
3.3.2.3 Special Narrowcast 11
3.3.2.4 Mobile Access 11
3.3.3 Constraints and Comments 11
3.4 WEB LINKS 12
3.4.1 Scope 12
3.4.2 Scenarios 12
3.4.2.1 Immediate Link Following 12
3.4.2.2 Automated Link Following 12
3.4.3 Constraints and Comments 12
3.5 SEGMENT JUMPING 13
3.5.1 Scope 13
3.5.2 Scenarios 13
3.5.2.1 Manual Segment Jumping 13
3.5.2.2 Automated Segment Jumping 13
3.5.3 Constraints and Comments 13
3.6 CONTENT CUSTOMIZATION 13
3.6.1 Scope 13
3.6.2 Scenarios 13
3.6.2.1 Locality Customization 13
3.6.2.2 Person Class Customization 14
3.6.2.3 Person Customization 14
3.6.2.4 Efficient Bandwidth Use 14
3.6.3 Constraints and Comments 14

4. REQUIREMENTS 14

4.1 CONTENT FULFILLMENT AND SELECTION 14
4.1.1 Conceptual Model for Content Identification and Acquisition 14
4.1.2 Selection 15
4.1.3 Fulfillment 15
4.1.4 Program Identification Requirements 16
4.2 CONTENT RIGHTS AND SECURITY 17
4.3 DEVICE AND CONTENT MANAGEMENT 18

4.3.1 Storage Devices *18*
4.3.2 Space Allocation *18*
4.3.3 Content Management *18*

5. IMPLEMENTATION CONSIDERATIONS 19

5.1 CONTENT USABILITY 19
5.2 RESOURCE USE 20
5.2.1 Resource Availability *20*
5.2.2 Resource Conflicts *21*
5.2.3 Communication with the User *21*

Table of Figures

FIGURE 1: SOURCE CONTENT IS LOCATED IN TIME AND SPACE 3
FIGURE 2: SELECTION PROCESS 3
FIGURE 3: ILLUSTRATION OF REQUIRED PROCESSES OF SELECTION, DECOMPOSITION AND RESOLUTION 5
FIGURE 4: CONTENT SELECTION AND FULFILLMENT PROCESSES OF BROADCAST CAPTURE 7
FIGURE 5: CONTENT SELECTION AND FULFILLMENT 16

Applications for Home

Storage and Internet Based Systems

TV Anytime and *TV Anywhere*

1. Introduction

The delivery of Audio-visual material in a digital form to the home creates the fundamental opportunity for consumer systems to store content on digital media, such as magnetic disks and tape or optical disk drives. This opportunity can be realized both in the PC and digital consumer device markets.

Storage technologies, driven by the computer industry, are cascading in price and capacity by factors that are almost unimaginable. In the last ten years there has been a one hundred-fold increase in the capacity/cost ratio of hard disk drives and currently the ratio is doubling every ten months. Incredibly, without any technological revolutions, this trend is reliably projected to continue, even to accelerate, for at least the next ten years.

In the year 2000, 10 GB of hard disk storage should retail at around $100, providing some four hours of audio-visual storage of MPEG-2 material at 5.5Mbit/s.

Assuming disk capacity/price doubling every 10 months and a more pessimistic 18 months, the following capacities at 5.5Mbit/s should be available for a retail cost of $100: -

Year	@ 18 months	@ 10 months
2000	4 hours	4 hours
2005	40 hours	240 hours
2010	400 hours	14,400 hours

Video content stored on disk for $100

Although today at a relatively high cost for the time stored in comparison with analogue tape, such disk storage will provide secure broadcast-quality storage, instant and random access, simultaneous record/play capability and will offer the opportunity for very simple control by users and agent technologies alike.

This document describes a number of applications which can be realized through the provision of home storage systems, made easy to use by appropriate use of content description, markers, links and agent technologies.

Material may be loaded onto a home storage device from broadcast content in real time, and also in non-real time, when the broadcast content is intended solely for recording. Material may also be loaded onto the home storage device from a remote storage system, in the manner of a *virtual video shop* using a video file transfer. The faster the connection, the more quickly the content will be available for consumption.

Disk storage will provide the ideal solution for time-shifting material, it will permit the viewer to start watching content before the recording has completed or to take a short break in the middle of live content without missing anything. It will also allow users to view material in a non-linear fashion, moving forwards and backwards easily with the ability to use defined index points and links where these are incorporated within content which has been so designed.

Tape storage systems may supplement disk, offering large volume random access storage. Also, there will be always be a requirement for content to be transferred from on-line storage to low-cost bulk media to be placed on a shelf for future consumption or to be passed to others.

The descriptive applications also consider the relationships between different media types, and the need, for example, for a storage device to manage a set of bookmarked WEB references obtained while consuming TV content.

Agent technologies will enable all of the above to be accomplished simply and easily; the agent making decisions about what should be recorded or deleted, managing the storage space according to the user's expressed wishes or previous behavior, based on information about the content being supplied or on offer.

The capability of IP based systems to carry audio-visual material in a digital form creates the opportunity to deliver AV material to any location which has a connection of sufficient capacity and quality. This enables users to access services and content from remote locations. Services considered here encompass television, and audio-only services from high-quality music through to low bit-rate speech.

2. Key Concepts

2.1 Source content is located in time and space

A fundamental assumption in the development of these applications is the temporal and spatial nature of content, which may be acquired by the local storage and network based system. Connectivity provides the means to obtain content which is not directly available at the point of consumption, and local storage provides the means to obtain content that is available in the future. Some material may be available both remotely **and** in the future.

In the diagram below the local storage system, or the user, views the world outside on two axes - one labeled *connectivity* and the other labeled *time*. Material that is accessible using an access network - as a file to transfer into the local storage system - is located directly on the axis of *connectivity*. Similarly, material that will be delivered via broadcast channels (terrestrial, cable or satellite) to the user's premises is located directly on the axis of *time*. Material which is located on a channel which is not accessible in the user's premises, but which must be accessed via a network is separated from the user in *connectivity* and *time*.

As time progresses, material on the right-hand side of the diagram moves towards the left, eventually reaching the *time zero* line. At this point material may be transferred into the local storage system, either directly or using the access network. It may also be transferred into a remote server, and become available for future transfer as a file.

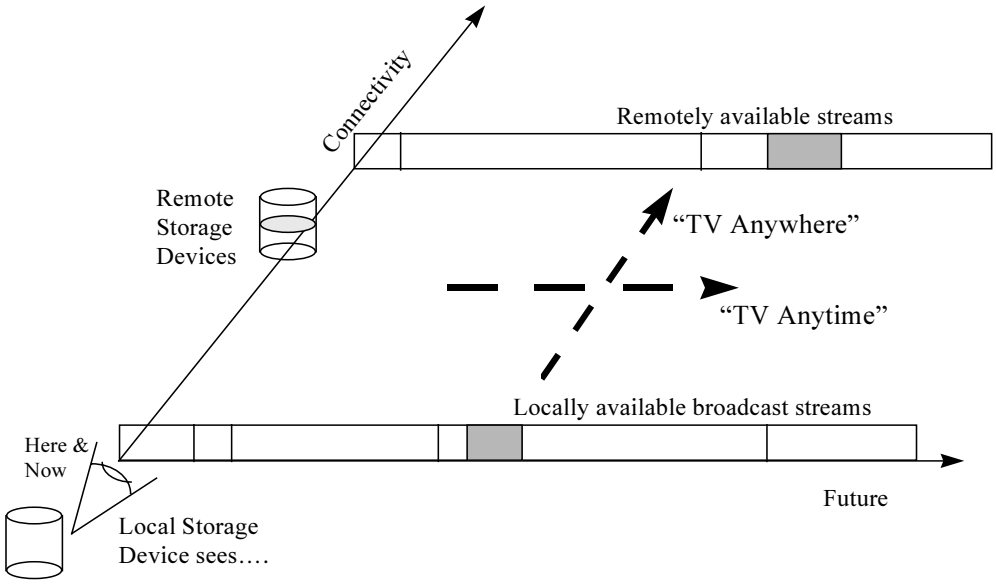

Figure 1: Source Content is Located in Time and Space

The diagram introduces the key concepts of *TV Anytime* and *TV Anywhere*, representing the temporal and spatial separation of a particular piece of desired content from the user's local storage system.

2.2 Selection

The logical selection of audiovisual material consists in the identification of material that meets criteria defined by a human. Selection can involve several steps.

In practical cases, several selections can be chained from general searches producing large sets through specific searches within sets to final direct selection from the remaining list of items. Selection can be user initiated or agent initiated. Selected programs can be either a specific program or a collection like a series or a thematic set.

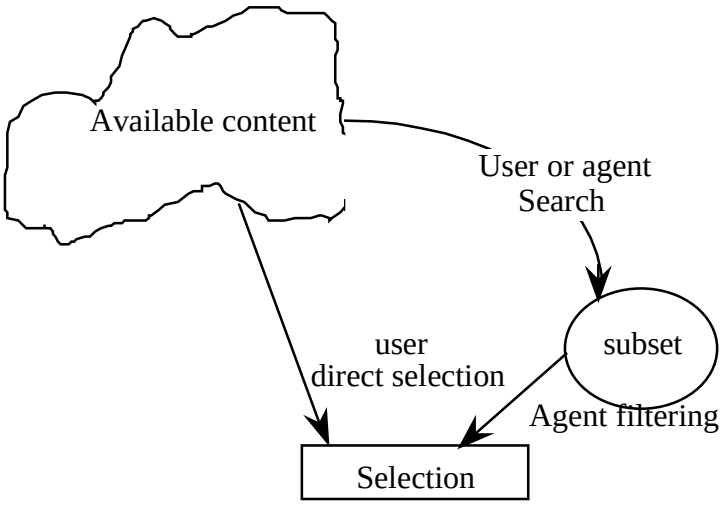

Figure 2: Selection Process

Selection mechanisms include:

- **Search-based selection**: A search engine or an agent accepts a request and produces a list of items supposed to fit the criteria. The search may use many data sources, including EPG data or program metadata.
- **Direct selection** (e.g. activates link on a web page or chooses from a list): Direct selection may be from an EPG, from a promotion, or from other printed or electronic sources.

There may be security issues related to the search and selection mechanisms, such as :

1) Closed group items. In some cases, only the members of a special group should know of a program's existance. Others should not have their search engines list this program.
2) Parental guidance and security ratings can be viewed as a special case of 1).
3) Security and/or cost metadata could be included as search criteria.

2.3 Fulfillment

Having selected and identified the desired content, the process of identifying the component elements, locating them in time and space, and their subsequent acquisition, is known as 'fulfillment'. The fulfillment process may be further broken down into three successive lower level processes:

- **Decomposition:** The selection process will have resulted in a unique human-relevant or computer ID, or collection of IDs, which may expand into multiple IDs. For example, a sequence or series of programs, such as the separate episodes of a drama, or a 'collection' of programs identified as being relevant to the user's selection criteria.
- **Resolution:** Having identified the elements, it will be necessary to determine where and when these may be found. A content element might be located in the local store, in a current 'live' broadcast, be the subject of a future broadcast, or available remotely on the Internet as a file or stream, either now or at a future date. It is necessary to keep track of any changes of location and/or timing of the required items; for example, a sports event taking longer than expected.
- **Acquisition:** Finally, having verified the user's rights to acquire the content, appropriate action is taken to obtain and store the content elements on the local storage device.

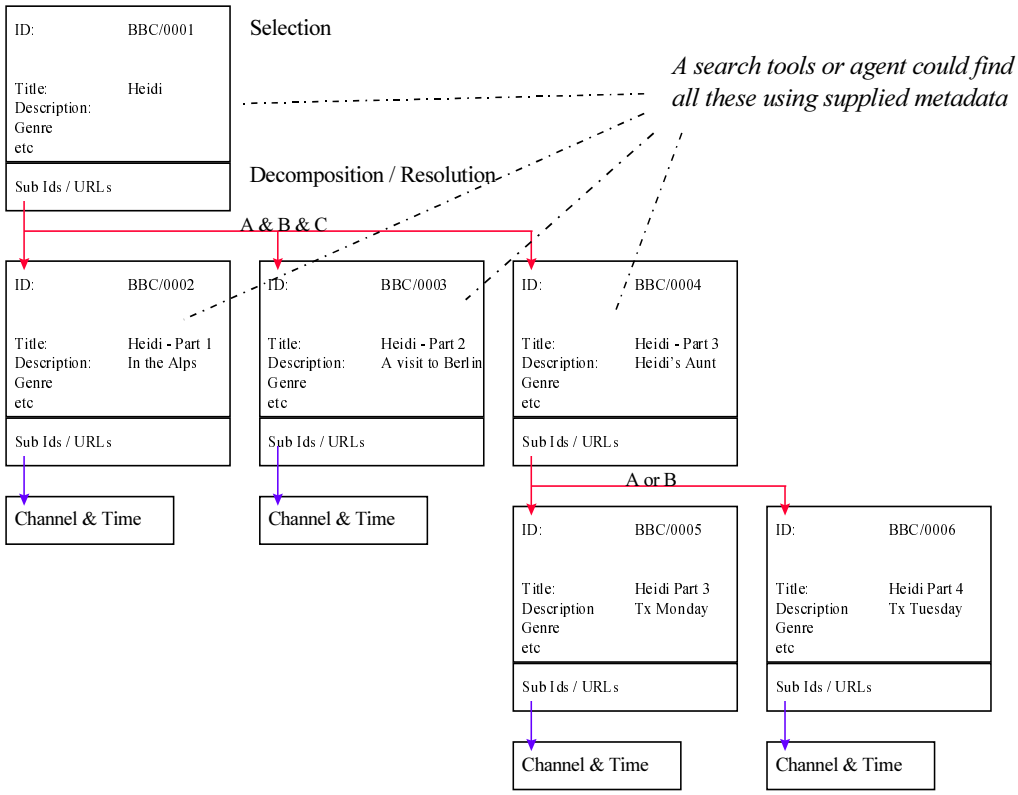

Figure 3: Illustration of Required Processes of Selection, Decomposition and Resolution

When a specific item of interest is selected, the system will be given a unique identifier that represents a persistent reference ID to the content, which is independent of time and space of delivery.

The resolving authority for this ID will then provide a mechanism by which the ID can be resolved automatically into a physical locator who specifies the time and place of broadcast. In some cases this process may require several steps. In the example Figure 3:, the series *Heidi* identified by *BBC/0001*. This is a collective ID, which decomposes into three episodes. When *BBC/0001* has been selected for recording, all three episodes will be captured automatically.

For each program item to be captured at least one physical locator is eventually required. In the case where multiple locators exist (e.g. repeated transmissions of the same content, where one locator or another may be used) the locator to be used may be chosen according to criteria such as time, cost, quality, or conflict with other content recording. In Figure 3:, Episode 3 of the *Heidi* series, identified by *BBC/0004*, resolves into *BBC/0005* or *BBC/0006*. Either may be used by the recording system.

While unique IDs can reference complex program structures based on relational links between content items, the search and filtering tools may deliver to the capture or fulfillment process a single ID corresponding to any point in the decomposition chain.

Once given to the capture process, fulfillment will result in the capture of all content items hierarchically decomposed from that point.

If a user or agent wishes to move up the decomposition chain – say where a user is attracted to capture a series having seen a single episode – mechanisms will need to be provided to allow users to derive any collective IDs from the resolving authority. Alternatively, the broadcaster may wish to insert explicit links to the series or any other related content into the content itself, using the same mechanism described for promos and trailers.

2.4 Management and Use of Acquired Content

TV Anytime and *TV Anywhere* will provide the most benefit if the local consumer device contains rewriteable storage for more than a single program. If the acquired material is stored, it may be viewed

at the convenience of the user. Depending on the rights conferred by the service provider, it may be viewed more than once or saved in longer term, off line storage. Most consumer devices with significant storage capacity will require one or more means for the user to selectively delete material.

When content is wholly or partially stored, then the content may contain structure to be easily exploited by the consumer device. A simple example is embedded links within the content to suggest alternate material of interest to the viewer. This could be references to material in local storage, or to material stored on publicly accessible networks. Such links would be invoked upon the explicit request of a viewer. It is essential that this process is graceful, easy, and intuitive for the user without the cognitive or interface complexity more typical of today's public networks.

A more complex example requires that content be developed with specific reference points within it and suitable linking mechanisms. The reference points, having been provided by the content producer, will surely be designed to be compatible with the content producer's commercial interests. One intent of such a structure permits the viewer to see, for example, the weather and economic news first and skip the sports. An elaborate example envisions content specifically developed to permit variety in the viewer's experience. This may include customized advertisements to be viewed on the basis of geographical location or viewer age, insert points to permit customized children's programming to the name(s) of the viewing children, or alternative content pre-transmitted for display at the appropriate time.

Management of consumer device storage space is important. The consumer device owner may wish to manage space among several family members or for different types of recorded material. Management may require policies for deleting material to make room for new content, for archiving material, for restricting viewing of selected material, etc. The general case of space management is very complex and may be difficult to implement while maintaining simple and intuitive interfaces for the user.

If the storage cost is high, a consumer device supplier may lower the consumer cost of a box in return for the right to use part of the local non-volatile storage. Example uses of such a right could include download of catalogs or web-like content with commercial benefits for the supplier. The storage rights granted to the supplier then must be ensured in competition with the needs and desires of the user.

3. Application Scenarios

A number of example application scenarios are developed in this document to describe the functions of a home storage based system, and indicate requirements of the underlying technologies, protocols and data elements.

The scenarios are classified into six categories.

- Broadcast Capture
- Video File Transfer
- Remote Stream Access
- Web Links
- Segment Jumping, and
- Content Customization

The first three categories include applications for locating and capturing content into local storage.

The last three categories include applications that can be performed after the content has been captured.

The presence of local (home) storage provides new opportunities for usage of audio-visual material in the home. Traditionally, this material would be viewed sequentially either live or from a recorded tape. If home storage includes devices that permit non-linear access to the content and include some degree of processing power, then many new forms of content use are possible.

Each scenario is reviewed briefly in the following sections:

3.1 Broadcast Capture

3.1.1 Scope

Broadcast capture applies to services and applications that are able to effect the automatic *recording* or storage of broadcast content for use by the user sometime in the future. The storage device may be a combination of long and short term, high and low latency media. Material that is to be captured may be transmitted in real time (that is at its presentation rate) or, when specifically intended for recording, in non-real time.

Broadcast capture may be:

- User initiated, where the user selects future content directly from an electronic program guide (EPG) or from a cross-reference contained in another piece of content (for example a promo).
- Agent initiated, where the user specifies certain characteristics in advance and the agent selects and records content that meets the characteristics.

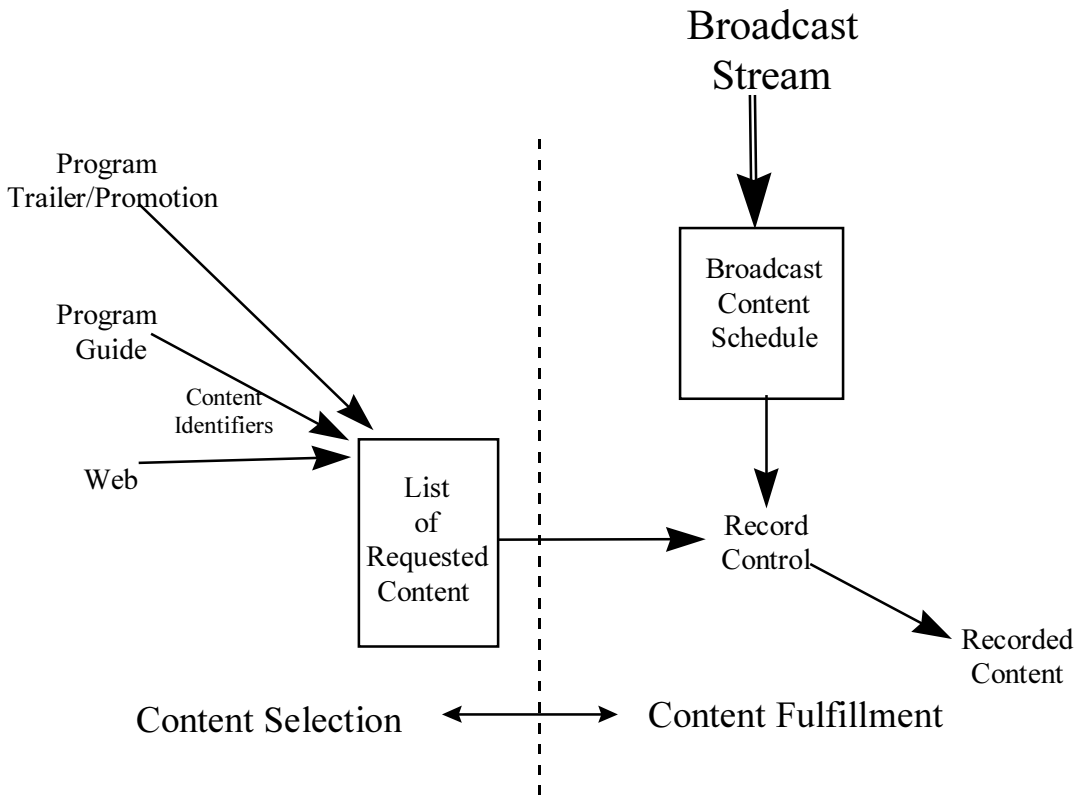

Figure 4: Content Selection and Fulfillment Processes of Broadcast Capture

In both cases content is selected according to information (including any information relating to payment) supplied in advance. Figure 4:illustrates how the selection process leads to the acquisition of content – the fulfillment process. Selection is the process by which the user or the agent selects specific content to record, based on a link obtained from a trailer or promo, through the user reviewing descriptive information via a program guide or Web, or through the agent matching the descriptive information to a set of selection criteria. Having been identified and selected, the content is listed for acquisition by the home storage system. The fulfillment process then requires mechanisms to enable the identified content to be acquired despite changes in transmission time and channel scheduling.

The home storage system needs to manage the disk space and to identify conflicts in content acquisition (when, for example, two pieces of content from different sources are scheduled for capture at the same time). More importantly, the system should be capable of operating in parallel with normal viewing and should be able to record content whilst a viewer watches any other material.

3.1.2 Scenarios

3.1.2.1 User Initiated Capture

Four possible scenarios are identified: broadcast capture using an EPG, broadcast capture using the internet, broadcast capture using embedded references in other broadcast content, and finally the facility of organized immediate recording including the capability to support dynamic viewing.

For broadcast capture via an EPG, the user is presented with an EPG containing, say, a list of scheduled content for the next week. The EPG will need to present data which is not only derived from factual and production information but will also requires a distinct marketing edge taking into account the user's perspective of the content. With the click of the remote, the user simply selects the name of the program that interests him or her. The application will then automatically register a unique identifier for that content and ensure that the content is recorded whenever and on whatever channel the program is broadcast, regardless of any subsequent changes in program scheduling. The application will also ensure that unused space of suitable size is found for the content or will inform the user of the need to delete some other content that is no longer of interest. After capture the user will be informed in order that the user may watch it at his or her leisure.

A similar scenario may be imagined for broadcast capture using the Internet. Here, the user may be attracted to a reference on a Web page of a program to be broadcast sometime in the future. The user simply clicks on the reference link and the program's id is automatically registered using the same capture procedure of the EPG scenario.

Another scenario uses embedded references in broadcast content. A program of interest may be brought to the user's attention during a broadcast. The user may then click on these references while continuing to watch the current program in order to automatically trigger the recording of the program of interest some later. An example might be during a trailer, when a user may register his or her wish to record the program without the need to know when or on what channel the program will be broadcast. Again, recording, storage and access procedures will be identical to that of the EPG scenario.

A service provider might be enabled to load content into the home storage device as a background task in non-real time without a specific request from a user. When the content has been loaded, the user could be informed that the title is available for viewing immediately, without the need to wait. After some period of time the content could be erased and replaced. The user could be charged for viewing the title only if it is actually watched.

Finally, there is the prospect of *immediate* recording. The user may wish to record a current program *on the spur of the moment* some time into the broadcast. Today, VCR technology requires that the user first find a free tape, establish that there is sufficient room available on the tape to hold the program, press the record button and hope for the best. With a home storage system implementing *immediate* recording, the user simply presses the record button. The system then finds a free area of storage and ensures that the program is correctly recorded – including any missed beginning portion. There may also be a facility to automatically detect the end of the current program in order to terminate the recording.

A further extension of *immediate* recording might include the prospect of dynamic viewing. For example, the user is watching a broadcast television program and wants to interrupt the program viewing. Then user may perform *Pause*, *Rewind*, and *slow-motion* type operations between the *near past* and the current broadcast point. While such an application may be less reliant on the program selection type attributes required to implement broadcast capture; it is assumed that Dynamic Viewing would utilize many of the same storage management procedures. In addition, Dynamic Viewing will require a simultaneous read-write capability

3.1.2.2 Agent Initiated Capture

Agent initiated capture can be seen as being a logical extension of the user-initiated capture scenarios where instead of the user explicitly selecting a program for recording, the agent automatically performs the task based a user profile. As such, many of the program attributes of user-initiated capture will be directly applicable to agent initiated capture, in addition to potentially many new attributes that will be required to profile different program categories.

Agent technology may require crisper categorization than is needed for user initiation. The categorization attributes of content will need to describe at the very least commonly used genres such as *news, movie, soap* or *documentary* and, in addition, to be expandable to include new genres such as *docu-soaps*.

Content can be classified according to multiple categories; e.g. the capturing of a broadcast item corresponding to a search of a *movie* of one particular subject matter would not preclude the same movie be captured as a result of specifying a completely different subject matter that is still relevant to that film.

In one example agent application, on installation of the system, the user is prompted to describe his/her personal likes and dislikes in TV viewing, based on genres, categories and types that are also attributes of digital TV content. This should be done in a manner that is pleasing to the user.

The user can specify

- programs to be captured on an ongoing basis, for *Anytime Services* (e.g. latest news, specific sports, soap installments, children's programs);
- features to be captured in case they come up, such as categories of movies (e.g. horror films, specific directors);
- programs featuring specific persons;
- programs featuring specific themes (e.g., crime, politics, technology).

The user also specifies personal dislikes, programs that should never be captured, e.g. pornography. Likes and dislikes can be given priorities, prompted by the system, based on storage capacity or other constraints.

The system then filters incoming program and service information, using the attributes from Service Information, information from specific web sites etc. Based on weighted filter output programs are automatically captured on local storage. Results of this capture process are presented to the user in a manner consistent with the local EPG. This presentation may include user-specified prioritization or ordering of the available live and recorded material. The user can select captured features for viewing in the same manner as selecting real-time broadcast programs.

The system can also be made adaptive by a second filtering process, taking into account the actual viewing habits of the user, again using the attributes assigned to digital TV program categories. The actually viewed categories may be different from the specified likes and dislikes. By suitable weighting of the specified and adaptive filtering processes the personal profile is adjusted, possibly with prompting from the system.

The system provides for automatic management of storage capacity, using specified priorities, prompting the user when necessary, e.g. to take a decision about what to delete for release of storage capacity.

The user should pay only for the services used. For example, a pay-per-view movie stored on a local storage device but not viewed should not lead to any payments.

3.1.3 Constraints and Comments

The descriptive data of the content needs to be defined and also needs to correlate with user expectation. Content should use globally unique identifier. There may need to be a defined interface to and from the storage device.

The availability and use of a set of standardized program attributes is a necessary interoperability requirement to make automatic capture work. These attributes must be provided by broadcasters and service providers in the MPEG-2 TS. Basic program attributes can be found in Service Information, possibly some extensions are needed. A future possibility to make the system more independent of this is feature extraction. Procedures and mechanisms will be required to maintain the relevancy of content attributes. In particular, a means must be found to deal with changing and evolving categories of content; e.g. content attributes will need to be extendable to include future new content attributes.

The system presumes a digital TV consumer device, a set of filtering mechanisms, local storage and an alluring user interface.

Automatic capture assumes some of the capture services tools offered by the broadcast capture scenario, the agent taking the place of the user.

3.2 File Transfer from Remote Storage Devices

3.2.1 Scope

This application provides the user with the capability to receive and store video files in non-real-time and to display them later.

3.2.2 Scenarios

In a *pull* file transfer the user selects a title using a graphical user interface from a list of available titles to be viewed at a later time. The waiting time before the title may be viewed is variable, ranging from say almost instantaneous access to a very slow download operation. High-speed delivery may be offered as an extra cost option or possibly in exchange for lower picture quality. When delivery of the title is complete the viewer is notified, either on-screen or by means of an indicator light. After delivery is complete, the title may be viewed as if it were on a videotape, including pause, rewind and fast forward functions. Different distribution and payment models should be supported, including:

- One time payment for unlimited use
- Period viewing
- Payment for a fixed number of viewings

Video file transfer may also be used in a *push* mode where a file is loaded onto the local disk without any action on the part of the user. Examples include video e-mail addressed to the user or movies loaded onto the disk during off-peak times. The user would be notified via an on-screen message that these had been loaded and invited to select and view them.

3.3 Remote Stream Access

3.3.1 Scope

The continuing evolution of telecommunications networks and IP-based connectivity can be expected in due course to lead to a dramatic increase in freedom for the user to gain access to material, irrespective of the physical or geographical location of either the user or the material. This will open up many more exciting possibilities for viewing and/or retrieving broadcast services and program content, unconstrained by the locations in which they are broadcast traditionally or stored.

The process of selection and fulfillment would be initiated either by direct user selection or by use of a search agent using the concepts described in Section 2. Content may be selected through a link on a web page or through a bookmarked favorite location. The user might have to select an appropriate level of service that corresponds to quality/cost requirements and available bandwidth. An IP connection would then be made from the user terminal equipment (maybe a PC) to the home service provider. Content may be watched in real time, or alternatively by retrieving it either in real time or in non-real time and storing locally.

3.3.2 Scenarios

3.3.2.1 Home Country from Hotel

A user on business abroad wishes to watch a home TV channel for an important event. To do so, the user selects a service, for example through a link on a WEB page or via a bookmarked favorite and an IP connection is made from the terminal equipment (for the traveler, usually a PC) to the home service provider. It may well be necessary to select between a number of possible services and to establish identity and/or location, in some way to gain access to the selected service. The establishment of user identity may require user authentication as well as protection of user privacy. The local terminal may also need to be authenticated and its location verified. The user will probably need to request a level of service that corresponds to quality/cost requirements and available bandwidth. The requested TV service is then streamed to the user's equipment at the requested or available quality. Using Home Storage-like concepts for caching may improve QoS. The equipment may need to resolve problems relating to different content formats (e.g. format and line standards).

3.3.2.2 Foreign Broadcast from Home

A user at home may wish to use an IP connection to watch TV services originating from far-distant networks. Motives may include the scheduling of special events at the user's place of birth, special events connected to the user's interests, etc. Accessing foreign broadcasts also enables people to venture out of their normal TV environment, participating in programs that are not available in their demographic local area, or to visit new regions through native TV programs. These services and applications are enabled by the connectivity offered by the global IP network.

Persons may view the content in real time, which may have to be done at abnormal hours in case of a large difference in time zone, or capture content on local storage. The terminal equipment can be a digital TV receiver, a STB or a PC. Selection of the programs uses the same user-initiated and/or agent initiated mechanisms as described under broadcast capture. Instead of a local EPG, broadcasters' or other web sites will help identify and locate programs.

Because the remote access could demand significant communication resources, this service may be appropriate for offering different pricing that trades viewer quality versus cost. The access, QoS and content format issues are similar to those mentioned in the Home Country from a Hotel scenario.

Material intended for this type of access may be designed with auxiliary languages to encourage viewers who do not speak the language of the content. It may also require novel arrangements to accommodate the interests of content rights holders, advertisers, local broadcasting regulators and others. (See Section 3.3.3 for more information.)

3.3.2.3 Special Narrowcast

A broadcaster may provide the service of making a specific program available on the Internet, say, after the actual broadcast has taken place. The availability may be more or less limited in time, according to the broadcaster's wishes and user demand. The storage function is provided at the service provider end.

3.3.2.4 Mobile Access

With the advent of IP networks and wireless communication systems it is now conceivable that A/V content (broadcast and non-broadcast content) be accessible from any location regardless of physical locations over time. In other words, content can be retrieved and viewed as the A/V receiving unit moves dynamically from location to location over time. A user of a mobile A/V system may wish to view their favored broadcast program while travelling to and from work or while on vacation. This would mean the broadcast content should be accessible via independent wireless networks. If the traveler ventures beyond the reaches of one wireless network the content should be seamlessly accessible from another wireless network. An example of this would be similar to cellular telephone networks.

3.3.3 Constraints and Comments

Almost all *TV Anywhere* scenarios require a re-assessment of the implications of contractual agreements between content owners, broadcasters, and viewers. Many of these agreements contain constraints on who may view, where, geographically, they are permitted to view, or when the viewing is permitted. Each of these constraints may requirement enforcement mechanisms in a *TV Anywhere* system or the contractual requirement may require modification to permit effective deployment of the service.

- **Personal Constraints:** Many countries restrict distribution of material by a person's age or social status. As an example, controversies over pornographic material commonly discuss the issue of whether the laws of the distribution location, the viewing location, or the laws of countries through which the material transits apply. Some countries restrict material of political or polemic nature. In most analog distribution infrastructures, the physical point of termination of a cable or transmit or receive antenna is sufficient to enforce such constraints. *TV Anywhere* technology may require means to determine who is requesting viewing, who is providing it, and maybe also the geographic location of end points or communication routes to enforce applicable regulations.
- **Geographic Constraints:** In some countries, viewers pay for the right to receive content and the content owners then restrict distribution to within country and to the licensed users. If such users are now in a different geographic territory, one can ask whether their original right to receive the content supersedes or is subordinate to the geographic distribution limitations of the content owner. A variation of this consideration can apply even if the

viewer has no purchased right to the material. An example is content that contains geographically distinct advertisements, possibly providing a commercial offer that is valid only in a limited territory. If such material is not permitted to be distributed outside this territory then either content must be customized for distribution or some access authorization mechanism is required.

- **Temporal Constraints:** Film distribution worldwide is conventionally timed differently at different locations. Recent controversies over DVD encoding techniques are evidence that preservation of this business model is important to the industry. If such material were now available via *TV Anywhere*, one must ask if the home (native) location of the person applies, or the viewing location, or the location of the distributor. Whatever the answer, means for determining the relevant location would be required or some current contractual distribution obligations would require revision.

TV Anywhere by its very nature must contain some discovery process to locate the requested material. If the experience of the Internet is applicable, then the search may locate several different instances of the same program (e.g. from mirror sites). In this case, selection of the site delivering maximum QoS would be desirable, but the choice may also be affected by applicable regulations as reviewed above.

If viewers come to see *TV Anywhere* as a logical equivalent to local TV service, this then suggests that the functions described in the broadcast capture scenarios of *TV Anytime* may be expected for *TV Anywhere*. Because many properties of *TV Anytime* are designed to support material with links to content available either on the broadcast or on locally accessible internet sites, this facility would be expected to work in a *TV Anywhere* situation.

3.4 Web Links

3.4.1 Scope

The Web Links application allows broadcast television content to provide the user with links to Web pages, to save links for future use, to manage a list of such links, enabling a DAVIC terminal or connected equipment to launch a WEB browser to access the related content.

3.4.2 Scenarios

3.4.2.1 Immediate Link Following

The user is watching a broadcast television program and a message appears on the screen indicating that additional information is available at a Web page. This message gives the user the options to either go to the Web site immediately (if the terminal is suitably equipped) or to store the link in a bookmark type of list for viewing later. The user selects the option to store it. The link is stored in a list and the message disappears from the screen.

3.4.2.2 Automated Link Following

The user has selected an option to automatically save all web links to the bookmarks list without asking the user. The user can watch the program undisturbed by on-screen messages. Afterwards the user can view the bookmarks list and use the saved web links. These may well be organized according to program boundaries and associated program text.

A second example might be a sports fan who chooses to always follow links to complete game statistics whenever available during sports telecasts. Such a link following could be automatic, with a return to the live broadcast whenever the user commands.

3.4.3 Constraints and Comments

The full capability Web Links application requires a digital broadcast television receiver that is capable of Internet connectivity and presenting Internet content. However, consumer devices that do not have any direct Internet connectivity will still benefit from the bookmark and list management facilities. In such a case, the user will need to be able to view the stored URL.

3.5 Segment jumping

3.5.1 Scope

Segment jumping is non-linear use of linear transmitted broadcast material locally stored on an in-home storage device. To use the material in this way it must be supplied with embedded links or markers such as index points and hyperlinks or a table of content.

3.5.2 Scenarios

3.5.2.1 Manual Segment Jumping

With content such as magazine shows, news or current affairs, content can contain links or markers that demarcate the natural components of the content. After the content is recorded on the home storage device, it is then possible to jump forward and backwards through the material using the links or markers. To facilitate this use, content may include a table of contents that the consumer device could display to facilitate use by the usual remote control

In addition to simple indices and links, content may also include *hot* links between items. For example, an educational program might display icons or text (e.g., in the form of a question) giving options to go to related topics within the content such that user's viewing perception of the content is that of an integrative experience. In this guise, the use of the broadcast content resembles the use of hyperlinked multi-media often sold today on CD-ROMs or used in electronic documents with embedded links.

3.5.2.2 Automated Segment Jumping

If content is marked and indexed for manual segment jumping, the facility can be useful both to the broadcaster and to the viewer. Segment jumping is one way to automate content customization by setting the consumer device to automatically follow some links. Such automated methods may include a user-specified sequence in which a nightly news story is viewed (e.g., weather first, followed by sports, national news, financial news, etc.) A broadcaster may also include instructions for segment jumping. For example, for ad customizations sets in a certain geographic area should follow one set of links, and sets in another area another.

3.5.3 Constraints and Comments

Segment jumping must have an acceptable default behavior if the local consumer device either does not implement the facility or if it has not recorded the content so link following is not feasible. The facility may wish to include paid links for subscribers, and free or clear links for non-subscribers. This means that access control would be required for components of content in addition to whatever is required for the overall content.

3.6 Content Customization

3.6.1 Scope

Content customization describes a means to present audio-video content with element insertion, deletion or adaptation that depends upon a local consumer device attribute and for which the inserted element may be retrieved from a home storage device.

3.6.2 Scenarios

There are many potential applications for customizing content at the home, rather than at the service provider. Such customization could include adaptations based on specific attributes of the viewer, or could permit the content to include material either pre-recorded or distributed previously to the home.

3.6.2.1 Locality Customization

A consumer device has its postal code or another locality attribute defined. Content with embedded alternative elements is broadcast in multiple versions on one or more program stream and the local consumer device is instructed to select those elements marked for the local attribute of the consumer

device. When the proper time occurs, the locally customized element substitutes for the generally available element in the broadcast stream. The capability could customize advertisements (e.g., to provide the local phone number of the nearest vendor of the product) or permit local or regional news highlight inserts (e.g., high school news coverage).

3.6.2.2 Person Class Customization

A consumer device has an attribute of the person watching. One use would be the age or gender of children that use the particular consumer device. Content with embedded alternative elements is broadcast in multiple versions on one or more program stream and the local consumer device instructed to select those elements marked for the local attribute of the consumer device. When the proper time occurs, the locally customized element substitutes for the generally available element in the broadcast stream. Program inserts might provide altered program ratings or ads customized to the gender, age or other preferences of person's watching.

3.6.2.3 Person Customization

A consumer device has stored the name of the person watching. Content with embedded alternative elements is broadcast in multiple versions on one or more program stream and the local consumer device instructed to adapt marked elements for the name of the viewer. An application of this capability could be program inserts to customize children's programming to the name of the child.

3.6.2.4 Efficient Bandwidth Use

Suppose a broadcast channel has bandwidth that is insufficient for some content. One example is high definition ads in a standard definition program stream, or the desire to transmit several different video renditions of an ad for display based on the persons preferences as noted in a previous scenario. To achieve this, the high bandwidth content would be delivered earlier than the presentation time, using available bandwidth, and stored for playback at full integrity at the proper time. The real-time programs would contain appropriate markers to trigger the insertion of the locally recorded material. A fully developed use of such a capability might transmit extensive material in late night hours for use during the following day period.

3.6.3 Constraints and Comments

Content must be developed to enable this behavior and the local consumer device requires a means of identifying the proper element(s) to record. To work properly in consumer devices without the capability the content must be designed to have integrity if the behavior is not implemented. In addition the local consumer device must have a means of identifying when and how to insert the locally recorded element. The local receiving device would require sufficient advance notice to accommodate the latency of A/V storage to ensure timely retrieval and insertion of content elements.

4. Requirements

4.1 Content Fulfillment and Selection

4.1.1 Conceptual Model for Content Identification and Acquisition

To obtain *TV Anytime* and *TV Anywhere* functionality, the user must identify A/V content of interest and the DAVIC system should then determine when and where that content may be obtained. The conceptual model of this process has a selection phase followed by a fulfillment phase. During selection, the user identifies one or more content items to be obtained. The fulfillment phase uses those identifications to obtain the content for the user. Differing implementations may hide or automate many required steps (emphasizing ease of use) or may reveal them or offer choices (emphasizing flexibility), but the basic strategy remains.

Every TV program contains parts. These parts could include leaders, dramatic segments, advertisements, etc. In the fully digital TV system, auxiliary parts (such as alternate languages) or related parts (such as Web page references) may also be present. To support *TV Anytime* and *TV Anywhere*, each part must have some *name* (i.e. unique computer processable ID) that is unique to the part. The material to which

the informal name *TV program* refers must have a unique name (i.e. unique computer processable ID) for the whole and may have unique names for many of its parts. Parts themselves may also have names (either a human relevant name or a (i.e. unique computer processable ID) for part components. Names (either a human relevant name or a (i.e. unique computer processable ID) may also refer to collections of programs. Unique names (i.e. unique computer processable ID) and name relationships (e.g., part-of, collective-name-for) are a foundational property of content that permits users to identify, locate, and access Audio/Visual material.

To permit a DAVIC system to identify material precisely, names (computer Ids) of collections of content, content, or content parts must be globally unique over a many year interval. Such names will inevitably be long, complex, and of no interest to users. In contrast, the names that people use may be ambiguous (e.g., there could be multiple broadcasts of the same material) or refer to a collection of material (e.g., the human relevant name of a dramatic program series). People may also name programs by their attributes (e.g. college football). Therefore, an essential property of a DAVIC *TV Anytime* or *TV Anywhere* system is the ability to translate a human-friendly ambiguous name (e.g. the 6 o'clock news) to a precise (computer ID) which can refer to a precise {channel, time, event id, transport id, etc.}. Such an expansion process may not proceed to completion when the request is initially made (e.g., a request for all future airings of a series). Thus the conceptual model does not assume that selection completes *in toto*, followed by fulfillment, but only that selection on an item-by-item basis completes before each item can be fulfilled.

4.1.2 Selection

The selection process begins when a person provides a human-sensible name or a name surrogate. The DAVIC system then expands these names into a collection of unique content names (computer Ids) that can be used for retrieval. The expansion process may be automatic or may require additional interaction with the person. Some systems may only use locally available EPG information, while others may use more extensive resources. Some selection methods may use attributes (such as genre) which may require access to metadata about content in addition to the name (either human relevant or computer ID) of the content. The selection process may be as simple as clicking on an icon during a promo describing the content available.

The collection of names that results may contain unique program references, or generic references to material that can only be expanded to a final, unique identifier in the future (e.g., an instruction to record all episodes of a dramatic series, some not yet be listed in the available EPG). The requirements below note particular types of useful naming methods that must be provided by a DAVIC system. Additional methods are not excluded.

Although the content naming system must accommodate names (either human relevant names or computer Ids) for programs parts (e.g. a news story within a news program), this does not require all broadcast material to so categorize or name material. A DAVIC system is therefore constrained by the naming practices of content providers and distributors.

4.1.3 Fulfillment

Once a unique name (computer ID) is available, the DAVIC system must determine how to acquire the designated material. This requires consultation with a *resolving authority* to translate the program or content name to (e.g.,) a channel/time for broadcast material or archive/name for a request for material. In typical use, the resolving authority may be the EPG and fulfillment may only require waiting until the proper time and then tuning to the proper channel. In other instances, fulfillment may require consultation with an external *resolving authority(s)* to translate the unique name(s) (computer Ids) to an archive and resource name for retrieval upon demand. It is possible that a name (either human relevant or computer id) may resolve to multiple instances where that material is located (e.g., repeat broadcasts). Fulfillment may also require additional consultation with the user. For example, resolved names (either human relevant or computer ID) may determine that a free, abbreviated version is available next week, or a paid, full version is available this evening.

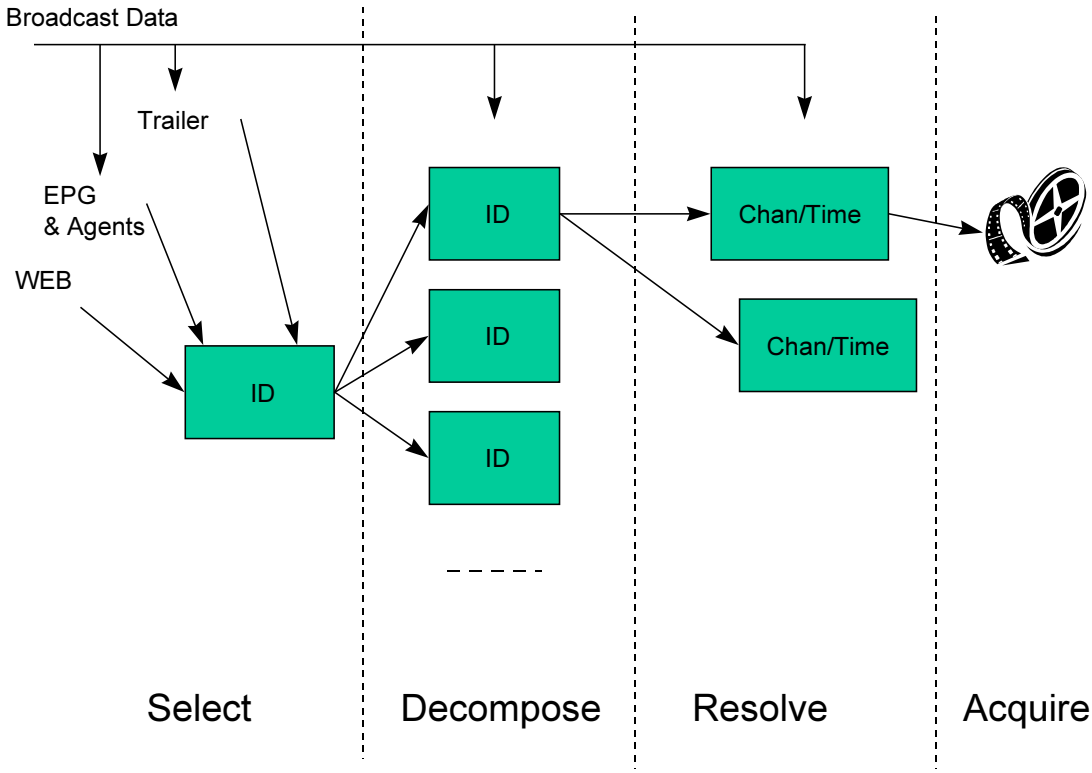

Figure 5: Content Selection and Fulfillment

Figure 5:illustrates how the selection process results in a unique ID which may expand into multiple ID's in the case of a sequence or series of programs which again resolve into locations in space and time where the program may be acquired.

4.1.4 Program Identification Requirements

A DAVIC system must provide the means to map program identification into unique content identification and the means to locate the material so identified.

To enable full use of *TV Anytime* or *TV Anywhere*, a DAVIC system must provide means for users to easily:

1) identify programs without requiring use of channel and time of broadcast. Program is the material implied by the public name of the content in a customary broadcasting interval (e.g., all material presented during the evening news program, including advertisements, leaders, etc.)

2) identify the channel and time on which a program is broadcast at specificity similar to current printed guides.

3) identify user-relevant segments of a program (e.g., news story, data component, advertisement)

4) identify sequences of program material (e.g., a TV series, news broadcast series) with a single reference. The requirement is intended to permit easy reference to multiple episodes of the same dramatic series, or to successive broadcasts of closely related content (e.g., news broadcasts).

5) identify multiple instances of material (e.g., multiple broadcasts) in which the potential differences in material are not of interest to the viewer (e.g. they include different advertisements).

6) identify instances of content related to but not identical with other instances. The system should permit users to select material that has user-valued dissimilar properties

(e.g., a full length vs. a movie edited for time constraints). Such content references may refer to entire programs or to program components.

7) identify program download opportunities & location if the system supports on-demand retrieval

8) identify legal or economic conditions on viewing or recording material.

If a system provides the capability to store unrequited material locally for viewing upon request, then the DAVIC system must provide a user-readable content description for the material. This requirement expects a description more informative than a brief title.

To support the capability for navigating within a program, a DAVIC system must provide means to

9) identify the current temporal position within a program

10) identify referenced content not part of the primary program (e.g., web-based material, stored material) consistent with equipment capability.

4.2 Content Rights and Security

The requirements for any security system deployed in a *TV Anytime* and *TV Anywhere* device will be defined according to the business demands of the content vendors and/or the broadcasters. In many cases, this will mean continued support of existing business paradigms (subscription, PPV) as well as future ones (rental, pay-per-time).

In order to facilitate a speedy development of *TV Anytime* and *TV Anywhere* systems, security mechanisms should initially focus on providing support for simpler business models, while permitting future extensions. Whatever system is chosen, it is inevitable that content protection cannot be absolutely unbreakable. Ultimately, motivated and well-equipped pirates will be able to obtain unauthorized access to content. What is important, however, is that their investments are rendered unprofitable. Piracy, with a significant initial investment will require the distribution of a large number of pirate modules (usually at a very low cost) in order to make a profit. Security tools, which prevent this type of piracy, are likely to be adequate for most cases.

“Free-to-air” content is generally broadcast unprotected. However, mechanisms that prevent unauthorized copying and distribution may still be required.

Taking all this into account, the DAVIC strategy for the protection of content in broadcast capture systems has therefore the following mutually supporting elements:

- Content authorizations, which must be clearly defined, communicated and accepted by the end user;
- An appropriate security system to prevent abuse of these authorizations.

The combination of content provider, service provider, security provider and network operator must have the ability to authorize users:

- to capture content on a local storage device, built into the STB or TV receiver or connected to a home network
- to use the captured material in a number of ways; e.g. one-time, many times, taking into account expiry dates etc.
- to transfer content to other media; e.g. with copy protection against unauthorized subsequent generations.

A DAVIC compliant end-user system must supply the means to enforce the content authorizations. In addition to being conformant with the content authorizations, the applications must be able to inform the user and the end-user system of the content authorization status. According to the business model, facilities such as free preview viewing should also be possible prior to the user being asked to pay. The application should be able to update the actual content authorization status after every use.

Security mechanisms should be put in place that protect the content across the appropriate external digital interfaces. Security levels will vary according to the opportunity for access to content via external interfaces. In the simplest case where a receiver and internal local storage together form a single integrated device without an external digital interface there may be no need for additional security

measures other than [tamper proofing and] those mechanisms already in place such as pay-per-view systems. At the other extreme, systems that support the transfer of content between local storage devices and removable storage may require the most demanding of security mechanisms.

Content transfer between local storage devices – via removable storage or the home network – should be transparent to the content security system in place; i.e. the act of copying content should leave the secure properties of the content intact and operational. One could even envisage the unrestricted distribution of content so long as the content provider can still guarantee continued revenue from pay-per-view systems.

4.3 Device and Content Management

4.3.1 Storage Devices

At the core of the system there shall be a random access re-writable storage system such as a hard disk. This shall allow concurrent random access read-write capability.

Additional storage devices may also be connected directly or via an in-home network. For example, tape, optical or other removable-media devices for archiving or for passing content to others. Some removable media may be read-only, such as DVD. Other media may allow read-write (such as DVD-RAM), or be write-once only (for example, some types of DVD-ROM).

Content may be stored off-line on removable media. The system would be aware of the presence of a removable-media storage device.

The content management system should identify the various storage devices available to the system, and whether media may be written to one or more times.

4.3.2 Space Allocation

The system must determine the available remaining recording storage space for each attached storage device.

There shall be support for the allocation of space according to criteria required to implement *TV Anytime* and *TV Anywhere* services. Allocations may need to be made for exclusive use by individual users or for direct use by one or more content providers. For example, the system might need to manage space among several family members or for different types of recorded material. Priority criteria would need to be set for each user in a multi-user system, perhaps including storage capacity limits for individuals.

There should be the ability to assign storage to one or more content providers to receive favorable payment terms in exchange for the right for the content provider to access the pre-determined storage area. In these cases, content may be loaded into the assigned area of the storage device in non-real time and without a specific request from the user.

The system should support a means of recording where new material automatically overwrites the old; for example, a weather report that is updated at regular intervals. It should also support a means of recording where explicit permission is required for the content to be deleted.

The system should record all relevant parts of the content, for example, video, audio, metadata and any other required data streams. The system should record appropriate information about content items due to be recorded at some future possibly not yet known date and take steps to initiate the recording at the appropriate time. It would be necessary to store associated metadata and attractors for each item of recorded program content. These would be used for search and retrieval within the local storage system, and also to avoid storing duplicate content. The user should be informed when requested content items have been stored.

The methods of storage should allow synchronization to be retained between each element of the content on playback.

The system must detect the end of a current program in order to terminate recording.

4.3.3 Content Management

A *TV Anytime* or *TV Anywhere* system will record content for later viewing. It must provide appropriate content management services for all material recorded within it or under its management control, including:

1) Suitable records of requests for content acquisition. These records need not be maintained after the acquisition is successfully completed;
2) Suitable records of content that was recorded by this device, including suitable media and location information appropriate to the content management strategy of the device;
3) Means to identify the nature of all recorded material to facilitate human assessment of the recorded inventory. The identification information will include suitable components of content metadata and appropriate information supplied by the device and /or user;
4) Appropriate catalogs or indices to facilitate human evaluation of recorded material;
5) Means to effect automated or human-assisted decisions for content archiving or deleting. This must include means to require human concurrence before content is deleted;
6) Means to support appropriate automated or human-assisted decisions for space management, including decisions for deleting or archiving of less valuable material;
7) Means to identify and maintain intra- and inter- content links in the material. This requirement applies only to the standard linking mechanisms appropriate to DAVIC content. It does not require the device to detect or interpret proprietary linking mechanisms. Also, it only requires that the system ensure that links to material that it manages can be effected appropriately. It does not require any assurance about links to material outside the management scope of the system;
8) Means to manage material recorded on removable media. The system must have a method to identify material previously recorded by that system;
9) If the content management system manages content on removable media, it must verify the identity and nature of content on media that it recorded when that media becomes accessible by the system;
10) If the content management system manages content on removable media, it must maintain suitable information in catalogs or indices to facilitate human evaluation and location of recorded content. The system must have a means to advise the user on appropriate identification for each removable media.

5. Implementation Considerations

Systems that support *TV Anytime* or *TV Anywhere* will make access to content appear easy and obvious to the user. Many challenges to achieve that goal demand sophisticated systems under the care and management of the broadcaster or transport provider. Customer equipment design is as challenging, particularly because the goal is to avoid conscious management effort and to conceal any difficulties from the user insofar as possible. The challenge for the customer equipment designer this is not to merely solve problems, but to do so in a manner that is natural and graceful to the user with minimal surprises and disappointments.

Two functions are particularly difficult to do well and are pervasive across a wide range of applications:

- Ensuring Content Usability
- Graceful Allocation of Bounded Resources

The quality of providing these two functions is difficult to measure objectively but failures are likely to be very evident to the user. This section reviews some of the issues for achieving the goals, not as to provide additional requirements, but to remind system designers of anticipated tensions between desirable system properties.

5.1 Content Usability

Content storage facilitates the delayed viewing of content and the viewing of the content in other than linear order. In addition, storage permits content to be acquired over an extended time and then viewed by linking the separate elements together automatically. To take advantage of these new possibilities, content may be designed to be a complex structure when it is delivered to the local storage, and this

structure may be essential for the proper viewing of the content. In addition, the content may depend upon material and resources external to the content itself that may not be available during delayed playing. Some examples illustrate the range of problems to be addressed.

- Content may contain internal references in the nomenclature of the transport technology that delivers the material. These references may need interpretation to be valid when encountered during playback.
- Constraints may be imposed on the original material that must be preserved at playback time. This may include the requirement for payment when the material is viewed (rather than at record time) or requirements for copy management or copy control.
- Content that is potentially customized at playing time (e.g. per person) may require that all alternative playback possibilities be recorded so that the customization can be enacted at playback time. Such customization may also require support from security resources that could find the proper authorization mechanisms unavailable at playback time.
- Content that has external references may not find them available at playback time. This may be particularly true for Web links that are timely such as weather or news.
- Content may be divided among many elementary streams. Suppose some of these streams carry alternative content, only one of which will be displayed depending upon some customization. Although only a relatively small fraction of the material may be required at play time, all of it should be recorded so the alternative choices at play time are feasible. This may require a much higher record rate than for a normal content stream.
- Content may reference pre-recorded content either delivered previously by the broadcast stream or via alternatives such as CD-ROMs. The recording process detects these references and takes appropriate measures to save the referenced material as well as the base content material.

The examples above illustrate the following types of challenge for delayed viewing of enriched digital content:

1) The content may have complex structure that must be preserved and processed during playback.
2) The content may have complex dependencies that are difficult or impossible to discern and preserve between record time and playback time.
3) The content may have constraints upon its use at record time that must be appropriately honored at playback time

5.2 Resource Use

These properties mean that *TV Anytime* and *TV Anywhere* systems must manage diverse requests for resources, determine appropriate action when resource requests exceed what is available, and communicate with the user when automatic resource management needs human assistance.

For example, *TV Anytime* requires the availability of local storage. The storage resource must be managed by the system, deciding what to record and what to discard. At this level, the problem is familiar to users used to VCR tape management. The new systems, however, will remember commands from users to record or process information at future or unknown times. Additionally, digital content is expected to sometimes include active behavior to execute either in the foreground or in the background while other viewing is in progress.

5.2.1 Resource Availability

A *TV Anytime* or *TV Anywhere* system contains multiple resources that must be shared. Example shared resources include storage space, available tuners, CPU cycles, the decryption/encryption system, communication equipment and displays. When a request is made to the system, it should determine whether the local resources are sufficient to honor the request before accepting it. This determination requires an accurate knowledge of resource availability when the request is executed and resources required by the request.

Resource availability limitations come from many causes.

- Permanent unavailability: this can occur by the limitations of the installed equipment or by changes or failures in the equipment.
- Conditional unavailability: this can occur because of changes in constraints imposed by the user, such as storage space allocation constraints, usage constraints during some hours, or blocking of some program types.
- Dynamic unavailability: this can occur because previous uses have claimed resources since the initial request was evaluated and not released them. Storage space claims are one example, as are tuner reservations for background tasks.

Because customer equipment will, most generally, support the viewing of multiple users, multiple claims for resources are inevitable. If these claims are for recording at a future time, and unless the resource is irrevocably reserved at the time of request, all requests may be honored. If in the course of honoring each successive request, the previous requests were not viewed and deleted, storage space exhaustion will occur even though each individual request found sufficient space when the request was made.

Not only is resource availability not always determinable when the request is first honored, it is also in general not possible to anticipate the resources required by the request until it is executed. For example, the time required to fully record a sporting event and the resulting storage space is not well known in advance. One approach to reducing this particular class of problems could be to require metadata that contains information on the probable minimum and maximum storage size of the content. A more comprehensive approach requires metadata to describe all embedded or active components that demand special resources. Even if such an approach were universally adopted, inaccuracies in the resource estimates and the variations in actual resource requirements (known after recording is complete) will continue to challenge efficient management of resources.

In summary, it is most appropriate to determine the feasibility of the request at the time of request and advise the user. If that is not possible, the system should determine feasibility before the request is to be executed. The least desirable behavior is to abandon a partially completed request or to achieve the request imperfectly. Thus the key issue for resource management is not merely to avoid deadlock over resource claims, but to provide graceful and appropriate behavior, hiding as much complexity as possible from the user.

5.2.2 Resource Conflicts

Although the best procedure is to not begin a request unless it can be completed, systems must also cope with unexpected conflicts over resources during the execution of a request. Partial equipment failure is a classic but infrequent cause. Other causes include storage space exhaustion by a background activity, unexpected claims on tuners by background activity or by user actions, unexpected processing claims by content that contains active material, or unexpected requests for communication resources caused by the action of the user.

As another example, it may be that a service provider subsidizes the cost of receiving equipment by claiming the right to record information in local storage. This could be used to provide speculative availability of paid programming, pre-provision of advertisements, etc. If such pre-authorized download were to begin when the user was otherwise using the tuner, decoder, etc. resources for viewing a different service provider, there is a substantial conflict over resource use.

As another example, suppose a receiver has two tuners. A previous request has claimed one tuner to record a movie while a user is watching another broadcast live using the second tuner. During that live broadcast, the user chooses to follow an embedded link that requires use of the second tuner. Does the system refuse to follow the link, terminate the recording, or request advice from the viewer?

Resource conflicts like these are particularly difficult because they often must be solved in a very brief time or the integrity of processing the content is compromised. The diversity of causes makes avoiding these problems very difficult.

5.2.3 Communication with the User

Part of the resource management processes may be sophisticated and well hidden from the user, but eventually problems require user attention. The challenge is to make the request for aid simple and easy to understand. Experience with personal computing equipment does not suggest this is easy to achieve.

In the general case, users should have some control over the algorithms to:

- State and modify resource use permissions
- Respond to resource reservation requests
- Arbitrate resource reservation deadlocks
- Identify and manage resource release

Even for knowledgeable computer system administrators, these questions are complex. Systems for home use must either radically simplify resource management at the probable cost of frequent inquiries to the user, or provide sophisticated resource management algorithms and then deal with the user's reaction to the inevitable puzzling behavior. Unfortunately, reacting to puzzling system behavior or to requests they do not understand, users often use radical solutions such as power cycling.

The goal for *TV Anytime* and *TV Anywhere* systems is a simple, easy to use, intuitive, and reliable device to provide easy access to content from any location and at any time. Although this section notes some difficulties to be overcome, other complex systems have achieved the goal. Telephone systems and automotive electronics are examples of complex systems that work reliably and unobtrusively. DAVIC *TV Anytime* and *TV Anywhere* systems should seek no less a goal.