

# An Architecture for Accessible and Sharable Digital Video in the Home Environment

Henrik Klinkmann  
Helsinki University of Technology  
henrik.klinkmann@iki.fi

## Abstract

Digital Video Recorders (DVRs) are becoming an increasingly important part of consumer electronics, especially with the introduction of digital television. We are seeing a revolution in the way televisions are utilized, but there are still stepping stones to be crossed on the road to a completely digital television system. Current DVRs work well and are rather user-friendly, but are especially lacking in that they do not allow easy access and sharing of the recordings stored on them. In order to acquire the user requirements for a system that would solve the issues of accessibility and sharability brief user interviews were arranged. Based on these requirements, we propose an architecture that solves these problems in a straightforward and easily implementable manner. Moreover, we propose how to implement this architecture.

**KEYWORDS:** DVR, digital video recorders, TiVo, Accessible and sharable digital video, television recording habits

## 1 Introduction

Television viewing has been an important element of family home activities since the nineteenfifties and -sixties. The videocassette recorder, or the VCR, came to revolutionize the way television programs were watched in the late seventies. Instead of having to watch a program at a certain time and place, it enabled the user to record and watch it at a later time. Video recorders kept evolving through the following decades, up until the late nineties. At this time, another invention came to revolutionize the way we watch television, namely the TiVo. The TiVo was the first mass-marketed product to find a large audience, and to introduce the endless possibilities of computer technology to video recording. It allowed timeshifting, skipping advertisements, advanced timed recording, a simple recording archive and plenty more with a simple-to-use graphical user interface [1].

Although newer versions of the TiVo now enable the user to plug in to and extend their TiVo with addons [2], a short check of the currently most popular digital video recorders (DVRs) for sale in Finland quickly reveals that DVRs are mostly meant to be stand-alone machines. The external connectors typically offered include audio and video output interfaces, and in some machines USB for file transfer to and from the hard disk mounted in the DVR.

What we call for is an integrated, extendable architecture for DVRs in the home environment. This architecture should

allow easy access and sharing of media files stored by the DVR; not only access from a television set, but from the wide ranging array of media electronics available in a home. This includes, at the very least, easy access to personal computers.

The focus in this paper is on the sharability and accessibility aspects of DVR media in the home environment. There are lots of other respects to take into account when discussing the recording, storage and sharing of video and other media files. Lots of research has gone into metadata related to the video files, and the secure and reliable storage of the files themselves. Directly related to these issues are the notions of filesystems and databases. We will touch upon all of these issues briefly and reference relevant research where suitable.

This paper is based on the results of a few user interviews that provide the basic user requirements of such a DVR architecture. Based on these brief interviews, we will suggest a simple architecture that solves the accessibility and sharability aspects. Furthermore, a possible implementation of the architecture is presented in the penultimate chapter.

A short glossary of relevant acronyms is found in the appendix. e

### 1.1 Related Work

Significant work is being done by several parties on the storage and archiving of memory artefacts, such as home videos, photos and more. Stevens et al. have produced an interesting article on an apparatus titled the Living Memory Box [5], which provides interesting insight into how people archive their memory artefacts, by building on ethnographic interviews and focus groups. Wactlar and Christer [4] have written a good overview on metadata related to digital video archives, listing all the relevant standards and processes and Abowd et al. [6] have researched metadata related more directly to home video archiving.

Storer et al. [7] have carried out an interesting survey on the secure and reliable storage of archived data, with a "write-once, read-maybe" methodology. Bolosky et al. [8] discuss the issue of utilising existing personal computers to build a distributed file system. Ficco [3] presents general guidelines for building a modern multimedia filesystem in his paper.

Q	User 1	User 2	User 3
1	23/f/student	22/f/student	25/m/student
2	two	one	two
3	digibox	no	no
4	45	700	300
5	no	VCR	VCR
6	TV	movies, TV, documentaries	TV, documentaries
7	two	one	two
8	no	no	once
9	never	documentaries once, movies 1-10, TV 1-2	once
10	weekends, vacations	evenings, when sick	evenings, weekends
11	two-three	15	none
12	no	movies	some documentaries
13	maybe	yes	yes

Table 1: Interview results

## 2 User Requirements

With the focus being on digital video in the home environment, we felt it was essential to get an understanding of what users request. The most direct way to collect user requirements is to interview a number of users. The target group of our interviews was young adults who are not that interested in how the technology works, only that it does work. This target group is interesting since they are likely to want to utilise modern technology, as long as it is easy to use and works well. The questions were oriented as to get the users current experiences of video recording in their homes, and also what they expect from an accessible and sharable digital video recording system for the home.

The aim was not to collect technical requirements for a potential DVR system, but rather to get an understanding of users' expectations and experiences of home recording and of their recording habits. With this knowledge, it is then easier to construct a system or architecture that solves the problems users currently are facing. It is important to remember that users tend to favor systems that do one thing well, rather than a hundred things half-heartedly.

### 2.1 User Interviews

Three supervised interviews were conducted in a relaxed environment with only the interviewer and the interview subject present. The interview subjects wrote their answers down themselves, and were told to ask if they had any questions regarding the interview questions.

The questions are presented in Appendix A. A brief summary of the answers is presented in Table 1. Note that the answers have been presented in a shortened form, to extract the relevant information.

### 2.2 Findings

The three interview subjects were all part of the same user group; that is university students, with a rather small personal interest in recording video. All three had rather extensive experience in the usage of digital TV set-top boxes (digiboxes), but only one of them owned one. None of the subjects owned a recording digibox (DVR), but two of them used a VCR for their recording needs instead.

Regarding the recording habits, there were slight differences between the users in that they all record (or would like to record) television series, whereas documentaries were more important for one user while movies were important for another. Typically, users watch recordings only once, after which the tapes are left to be rerecorded. One user likes to store certain movies to watch them up to 15 times, while another really does not keep a store of recordings. All users watch recordings in the evenings or in the weekends; one user mentioned being more willing to watch recordings when sick. All users would like to be able to access their recordings in other ways than via a television set. This is understandable, since all of them were active computer users, and owned at least one personal computer.

The interview questions were designed to acquire the user backgrounds on personal recording and requirements for a sharable and accessible recording system without asking that out front. The feedback received after the interviews was mostly positive, and there was great general interest in a more accessible digital recording system. However, one user complained that the questions were too complex, and that the results would have benefited from having more straightforward questions.

What can be said about this specific group of users is that they are very sparse recorders, and are not used to recording that much. This can be at least partly accredited to the fact that recording and watching recordings on a VCR is a rather tedious process. Other experiences from users who have started using the TiVo have noticed a large change in recording and television watching habits [1]. Television watching turns from a traditional broadcaster-controlled experience, to a user-controlled experience. Since none of the interview subjects belonged to this group, this cannot be confirmed nor discredited.

In essence, what the users request is a system that would allow them to easily record TV-series, movies and documentaries and to be able to access these not only from the television sets, but also from their computers. Since low cost is also something that users deem a rather important requirement, the system would need to be able to utilise existing technology to the greatest extent possible. The storage capabilities of the system need not be huge, but large enough to hold at least 40 hours of recordings (based on the storage patterns found in the user interviews), or the equivalent of 15 long motion pictures. Assuming the recording habits would be affected by the introduction of such a system, more storage space is highly likely to be necessary.

It would be very interesting to arrange more interviews, especially with a larger diversity of user groups. Interviews with heavy users of the current batch of DVRs would provide great insight into how they feel the DVR affects their recording behaviour, and what restrictions and problems users find

with these machines. However, already the small user survey performed here is enough to proceed to the architectural design of a recording system, which is not only easily accessible and highly sharable, but also easy-to-use and cost-efficient.

### 3 Architectural Design

With the findings from the user interviews as our most important motivators, we set out to design the blueprints of a system that would solve the main requirements. We did not restrict ourselves to what the users requested, since we wanted to design a system that is both feasible and efficient, technology-wise. Another major restraint on the design is the issue of cost. The resulting apparatus has to be within the economic boundaries of a normal Finnish household. We decided for the cost roof to be 500 euros, or the price of a current state-of-the-art digibox. Since the design is more of a theoretical nature, the price roof is not strictly enforced, but rather something that guides the design process.

#### 3.1 Existing Environment

The designed system would not be entered into a void, but has to be able to utilise and connect to the existing environment in which it will exist. This means that it will have to be able to interface with the existing technology of a typical household, as well as adapt well to its physical surroundings. What does a typical Finnish family household contain then?

That of course varies greatly, but we have decided to focus on the user group which was interviewed; that is young university students living either alone or with a partner. The people we interviewed lived in households with an analog, wide-screen television, a VCR, possibly a digibox and one or two computers, running Windows XP, and with Ethernet or WLAN network interfaces. These are all devices with which we intend to connect, in one way or another. Naturally, the system has to be able to decode the television signals broadcasted over the DVB system.

The architecturally most significant decision that has to be made is whether to incorporate currently existing digiboxes into the architecture through the use of extensions, or to design an alternative way to handle the task of decoding, recording and presenting the digital video content, or in other words to replace the digibox. Both methods are possible, but each has its own advantages and disadvantages. We will attempt to analyze and compare both options in the following chapters.

#### 3.2 Access and Sharing

While current DVRs do have a lot of features and bells and whistles, most of them are still lacking in the areas of access and sharing of their contents. They are simply storing and reading recordings onto their internal harddisk, with no standardized control-capable interface to the outside. Some DVRs provide a USB-connector, but the usage of the interface is limited to the sharing of media files. It provides no possibility to control the digibox externally. The TiVo, which is only available in the United States and Canada, has

more extensive capabilities, but is still not sufficiently user-controllable [2].

The user has to be able to access the media on the DVR from the outside, i.e. share it to the rest of his home, at the same time the DVR has to be able to read in video from the outside, i.e. provide access to the outside. A suitable metaphor for this is an island of media compared to a shared network of isles of media. The way to achieve a shared network of isles is to connect to the existing infrastructure in homes. In essence, this means to add connectivity to the DVR to a pre-existing Ethernet network, or to a WLAN, or even preferably to both. This would allow for easy access and control to both the contents of the DVR and for the sharing of contents from external sources to the DVR. Exactly what control interface the DVR provides to the outside is intentionally left undefined at this stage. What it needs to provide is control of the recording interface, so that recordings can be set up and changed externally, and access to view, modify and augment the video data and metadata stored on the DVR. It does not need to provide full control of the inner workings of the DVR, but it has to be possible to add, delete and access the recordings on the DVR, as well as read and modify the metadata contents and descriptors.

#### 3.3 Storage and Data

Digital television recordings consume a lot of disk space, creating a huge demand for secure and reliable but also inexpensive storage. With HDTV paving its way, the storage space requirement will eventually be even more important.

There are many possible scenarios how the storage system can be set up. Since we are looking for a simple and efficient system, we are not even going to contemplate fancy setups, such as distributed filesystems [8], although they would be possible. The model of storage [7] aimed for is a write-once, read-many, where the stored recording will be written to disk once, after which it has to stay consistent and readily accessible. In the interviews, users presented a need for storing some recordings for a longer time, which requires a sufficient level of reliability. It is hard to state what a sufficient level of reliability is, but given the small physical size and large logical size of hard disks, there is a significant risk in storing everything on a single hard disk only, compared to VHS tapes or DVD discs, which are separate physical entities. We cannot consider solutions where old data is archived on some safe, but slow-access media. Nor could we consider a fast, secure and reliable storage solution, since that would be inherently very expensive. Instead, we opt for the KISS-principle, i.e., Keep It Simple, Stupid.

The solution is then to use normal consumer ATA disk drives, but instead of using only one drive, we propose a simple RAID-setup [11] with two (or more) drives. Disks are prone to errors, often regardless of external factors which are accounted for disk failures [9], making data redundancy a viable solution. The compromise between reliability and cost, which is unavoidable, is solved by using a RAID-1 setup, where two disk drives mirror each others' contents. While there is a risk that one disk fails, it is improbable (but less likely) that both disks fail at the same time. With this setup, the system will not lose data in case only one disk fails, under

most circumstances. This setup will also not cause a noticeable performance hit. In fact, the read performance will be slightly better than using only one disk, but the write performance will be slightly worse. However, as the storage write model is write-once, this should not really matter that much.

Another important aspect to take into account on the subject of storage, is the question of what is to be stored, and how. In other words, in what formats will data be stored, and what is the storage hierarchy to be used. As an example, the MyLifeBits project [10] started out with a directory based, very hierarchical file system but in the end ended up with a relational database and a flat storage solution. This worked only thanks to the verbose annotation, or metadata, used in the project. There exist numerous standards for metadata describing video content [4], MPEG-7 [12] and MPEG-21 [13] being the most standardised of the bunch. While there are clear benefits with using a complete, standardised annotation system, these are very likely to be overkill for this application. Instead, an annotation method where the most important metadata properties are included, such as time and date of recording, TV channel, content title and length, content type and copyright information is likely to be much more usable. Imposing a huge standard on users, the likely result is that it will not be used at all or heavily misused [6]. We think an annotation system where the above properties are recorded, possibly using a subset of a complete metadata standard, such as MPEG-7, is a better solution. For metadata to be useful, it is essential that as much as possible and as accurate data as possible is recorded, which is why automatic metadata recording is a must. The system should be able to acquire most of the aforementioned annotation properties automatically, either from the television broadcast stream, or from an electronic program guide (EPG). The user can then manually change or augment the metadata, at will.

The video and audio data are proposed to be stored in MPEG-II format, most notably for their industry-wide acceptance and use, but also because of the fact that it is an open standard. Other popular media formats should also be supported, but the recording would default to storing MPEG-II formatted data.

### 3.4 Security and Usability

Security in the context of digital video in a home environment is a wide concept. However, the most important elements in the context of this document is the fact that intruders are unable to access the DVR home intranetwork. It also means that users of the system are not able to, even accidentally, share contents of the system to the outside. The reasons for this are mostly legal, but since there are huge legal implications involved, it is safest not to allow connectivity to the Internet.

On another note, the user interviews confirmed that usability is a key factor in the adoption of any new technical device. Therefore, a simple and clear user interface has to be a central part of the DVR. New users should be able to pick up the remote and just start using it. This was one of the biggest contributors to the success of the TiVo. However, more functionality usually means a more complicated user interface. The notions of accessibility and sharability

complicate things. One possible solution to this conundrum is to provide different user interfaces to users with different needs. Still, the typical user experience should be the focal point, offering both a simple interface and the main functionality. Details of the user interface are not handled here, but should be handled in greater extent, if this system were to be implemented.

## 4 Implementation

This chapter takes the results from the architectural design into consideration and suggests one possible theoretical implementation, with as many of the user requirements as possible fulfilled. We set out with designing a system that is capable of recording digital television in a manner that is easily accessible and sharable in a home environment. The system should be technically implementable, inexpensive and currently possible. These are not simple requirements; a lot of them are in conflict with each other. Thus, we have separated the implementation into hardware and software parts, for clarity's sake. Since the software is largely dependent on the hardware, that is handled first.

### 4.1 Hardware

In the architectural design it ultimately boiled down to two options; either extending on an existing DVR or building a new one. Since the first option is technically not viable currently, due to the lack of control interfaces on DVRs, the only choice left is to build a new one. This can be achieved by using a normal PC that has been designed with some special considerations in mind, such as size, noise level, extendability and standards-compliance. This is what is commonly referred to as a HTPC, or more precisely a HTPC DVR.

Since it is important to keep the design in check with reality, in other words the cost level, we present a short cost calculation of one possible HTPC implementation in table 2. The price of 545€ did exceed our price roof of 500€, but on the other hand this example was not the most bare-bone solution conceivable. The motherboard includes integrated video, sound, networking and RAID capabilities. The processor and memory exceeds what is minimally required, which is also true for the storage space. It should also be noted that this computer could well be used for many other things, including being a home server, a DVD player or even a gaming machine. This example HTPC DVR is lacking a composite video out connector, which can be added as an optional component, if the user has this need.

### 4.2 Software

The PC that was assembled in the previous section will not accomplish very much without software. Open standards and simple usage were noted as important requirements, thus the suggestion we make software-wise is to use Linux as the operating system, it being both open and widely in use. To provide the user with DVR capabilities there are several viable options, but the one that is perhaps most widely used in the HTPC community is MythTV [14]. It provides typical DVR capabilities, such as decoding DVB broadcasts, EPG

Component	Price (€)
HTPC mATX casing	90
mATX motherboard	80
Processor	65
Main memory	55
2 * 250 GB IDE HDD	165
DVB PCI card	70
Remote control	20
<b>Total</b>	<b>545</b>

Table 2: Hardware implementation

downloads, television viewing functionality and recording capabilities. It utilises a MySQL relational database for its metadata storage.

Since MythTV is GPL software, it can be easily extended and modified for our needs. The main changes that we would need to make are related to the metadata and to the accessibility and sharing capabilities. Most of this can be achieved using readily available tools and libraries for Linux. The metadata storage model would need to be refined, and an automatic recording of the metadata is a high priority. Furthermore, the possibility for the user to annotate the recordings himself is also an important issue that needs addressing.

### 4.3 The Middleware

To tackle the issues we are most interested in, we will continue the analysis of the problem domain. The file storage solution itself best be a flat storage system. The needed hierarchy would be provided via the metadata introduced from using the relational annotation database. In order to be able to access all this from the outside, i.e. from other devices, we need an access method to both the database and the file system. The file sharing could easily be solved through the use of Samba [15], a SMB/CIFS implementation for Linux (and other platforms), which allows transparent reading and writing of files over any underlying transport network. Further, since the MySQL database can be connected to using a normal TCP connection, that access problem is easily solved as well.

As for the easy usability, we cannot and do not want to rely on the user to manually connect to the file system and the database. To address this, we need software for the client devices connecting to the HTPC DVR. The client software should integrate into its host operating system in a way that it can be used as a middleware to provide further extensions onto this sharing architecture. It would allow for the simple reading and searching of metadata, as well as for the sharing of the actual video content, but also for the addition and changing of content on the DVR itself. It would also provide the ability to share content from client PCs onto the DVR as transparently and easily as possible. Of course, in the name of openness this middleware would be implemented in a cross-platform manner and published with an open license. The middleware would provide a viable platform for further extensions, in the form of, for example, client applications provided by the community at large. Ideally, this middleware architecture would provide a working example

implementation for an industry standard for media-sharing DVR architectures. The protocol that is used between the DVR and connecting devices is left an open question which will not be addressed here, but this protocol would definitely have to be open and preferably standardized. Exactly how the middleware would be integrated with MythTV is another important question left intentionally unanswered, however it would have to be built on top of the control protocol discussed above.

Another, rather different, solution to the client software, would be to provide a web service from the DVR, that could be easily accessed from any client device with a web browser. While this would provide better accessibility, it would slightly limit the usability and extendability of the overall architecture. Also, a web service could easily be integrated as an add-on to the existing architecture that would provide a simple, easy-to-use interface, but for more demanding users, we see the middleware approach as a more viable solution.

## 5 Conclusion

This document has made an attempt to achieve an open, simple, cheap, currently implementable architecture for a DVR system for the home environment with easy sharability and high accessibility. What has been achieved?

Basic user requirements of one specific user group has been documented through the use of a small survey. While it did not answer all questions regarding the requirements, it did give a taste for what users want. This was exactly the goal with the interviews. Naturally, in case someone really wanted to find out all, or most of, the user requirements, a bigger group of subjects would have to be interviewed. We would be especially interested in hearing the opinions from the group of heavy users of current DVRs.

With the basic set of user requirements acquired, the general outline of the architecture was drawn, with mainly the access, sharing and implementability aspects in mind. It was quickly derived that if this architecture were to be readily implementable, one would have to forgo of current DVRs, since there is no easy, standardized way to extend on them. Furthermore, other important aspects of DVRs were also taken into account. Safe and reliable storage were key elements in the concoction, as well as were safety and usability.

Lastly, a theoretical, but possible, implementation was suggested. It builds on a normal but specialized PC, using specially designed, open, software. The most important part of the puzzle to solve was the access and sharing aspects. These were handled through middleware that was to run on the client computers, connecting to the DVR. Again, openness and extendability were key factors.

To conclude, there is a greater need now than ever for a simple, efficient and inexpensive digital recording system. The current market is based on unstandardized, closed systems, which although provide lots of functionality are especially lacking in terms of access. Through the use of existing technology and networking, we have arrived at an architecture that allows users to forgo from current television and recording habits, and start being able to control their own recording and viewing habits.

## A Glossary

Word or acronym	Meaning
digibox	digital TV set-top (receiver) box
DVB	digital video broadcasting
DVR	digital video recorder
EPG	electronic program guide
GPL	GNU General Public License
HDTV	high-definition television
HTPC	home theater PC
MySQL	a free for personal use database management system
RAID	redundant array of inexpensive disks
TiVo	a popular brand of DVRs in the United States
VCR	videocassette recorder

Table 3: Glossary

## B Interview questions

- Age, sex, education?
- How many persons live in your household?
- Do you have a digibox/DVR/TiVo in your household?
- For how long have you used a digibox/DVR/TiVo (yours, or somebody elses) (in days, approximate)?
- Do you record on your digibox or Videocassette Recorder (VCR)?
- What kind of material do you usually record (or would like to record) (e.g. TV-series, movies, documentaries or any other material) on your digibox or VCR?
- How many personal computers do you have in your household?
- Do you have an archive of personal video recordings (family recordings, etc)? How many hours does it comprise? (approximately)
- How many times do you usually watch recordings you have made on your VCR or digibox? What kind of material do you watch mostly? Please be as detailed as possible, since this of course depends on the kind of material recorded.
- When (time of day and week) and where (what place, with whom) do you usually watch recordings you have made on your VCR or digibox?
- How many recordings (personal or VCR/digibox) have you lost, either misplaced, due to own mistake or wear and tear?
- Would you like to keep your recordings usable for a long period of time (over five years), or is it OK if they are lost? Again, try to explain what kind of recordings this applies to.
- Would you like to be able to access your recordings in other ways than via the digibox? For example, from your personal computer.

## References

- Roger Allan. The TiVo Box Redefines Television Viewing...As It Creates a New TV Viewer Lifestyle. In *Electronic Design*, February 2003, pp. 41-47
- Arthur van Hoff and Adam Doppelt. Building on TiVo. In *Dr. Dobbs Journal*, March 2005, pp. 64-70
- Mike Ficco. How to build a multimedia filesystem. In *Embedded Systems Programming*, March 2003, pp. 14-15, 18-23
- Howard D. Dactlar and Michael G. Christel. Digital Video Archives: Managing Through Metadata. In *Building a National Strategy for Digital Preservation: Issues in Digital Media Archiving*, April 2002, pp. 84-89
- Molly M. Stevens, Gregory D. Abowd, Khai N. Truong and Florian Vollmer. Getting into the Living Memory Box: Family archives & holistic design. In *Personal and Ubiquitous Computing*, July 2003, pp. 210-216
- Gregory D. Abowd, Matthias Gauger and Andreas Lachenmann. The Family Video Archive: An annotation and browsing. environment for home movies In *Proceedings of the 5th ACM SIGMM international workshop on Multimedia information retrieval*, 2003, pp. 1-8
- Mark W. Storer, Kevin Greenan and Ethan L. Miller. Long-Term Threats to Secure Archives. In *Proceedings of the second ACM workshop on Storage security and survivability*, 2006, pp. 9-16
- William J. Bolosky, John R. Douceur, David Ely and Martin Theimer. Feasibility of a Serverless Distributed File System Deployed on an Existing Set of Desktop PCs. In *ACM Sigmetrics Performance Evaluation Review*, June 2000, pp. 34-43
- Eduardo Pinheiro, Wolf-Dietrich Weber and Luiz André Barroso. Failure Trends in a Large Disk Drive Population. In *the Proceedings of the 5th USENIX Conference on File and Storage Technologies*, February 2007, pp. 1-12
- Jim Gemmell, Gordon Bell and Roger Lueder. MyLifeBits: A Personal Database for Everything. In *Communications of the ACM*, vol. 49, Issue 1, January 2006, pp. 88-95

- 
- [11] David A. Patterson, Garth Gibson and Randy H. Katz. A Case for Redundant Arrays of Inexpensive Disks (RAID). In Proceedings of SIGMOD Conference 1988, June 1988, pp. 109-116
- [12] Martínez, J.M. Coding of Moving Pictures and Audio, MPEG-7 Overview. In ISO/IEC JTC1/SC29/WG11N4980. URL: <http://www.chiariglione.org/mpeg/standards/mpeg-7/mpeg-7.htm>
- [13] Jan Bormans and Keith Hill MPEG-21 Overview v.5. In ISO/IEC JTC1/SC29/WG11/N5231. URL: <http://www.chiariglione.org/mpeg/standards/mpeg-21/mpeg-21.htm>
- [14] Isaac Richards et al. MythTV. <http://www.mythtv.org>
- [15] The Samba team. Samba. <http://www.samba.org>