

CatDV Pro Workgroup Server

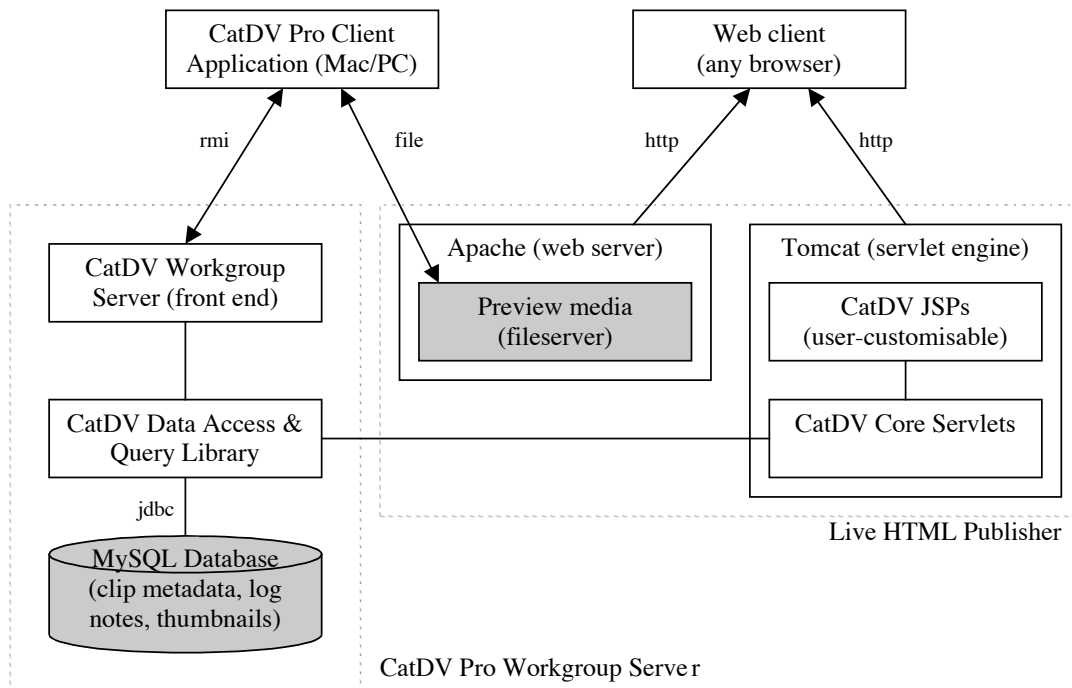
Square Box Systems Ltd
May 2003

Architectural Overview

The CatDV Pro *client application* is a standalone desktop application, providing video logging and media cataloging capability and available for both Macintosh and Windows PCs. It is based around a catalog or database of clips, where each clip represents either a media file (identified by filename) or a scene within a movie file or video tape (identified by tape name and timecode values). Each clip is annotated with a rich collection of metadata, including name, log notes, date of recording, a thumbnail image, video and audio format, and so on, which can be used when searching for clips or to provide a record of the contents of a tape even if it's not currently online.

In addition to its well-established standalone mode of operation (whereby files and catalogs are stored on the local file system), the CatDV Pro application can also operate as a network client to an optional CatDV Pro Workgroup Server.

The *workgroup server* is a separate program, running on a server machine and accessible over a local area network from one or more client machines. The server stores the clip metadata in a relational database, rather than a local flat file. This allows the clip database to be searched and updated by different clients concurrently, and permits much more sophisticated searches to take place.



Using a CatDV Pro client communicating with the workgroup server, users can search for existing clips in the database, update these clips (eg. by adding further log notes) and create new clips in the database (eg. by scanning a new tape, importing new media files, or by creating secondary clips from existing clips). Exactly the same operations, with exactly the same user interface, can be performed as when CatDV Pro is being used as a standalone application, the only difference being that the clips are saved and loaded over the network into a shared database.

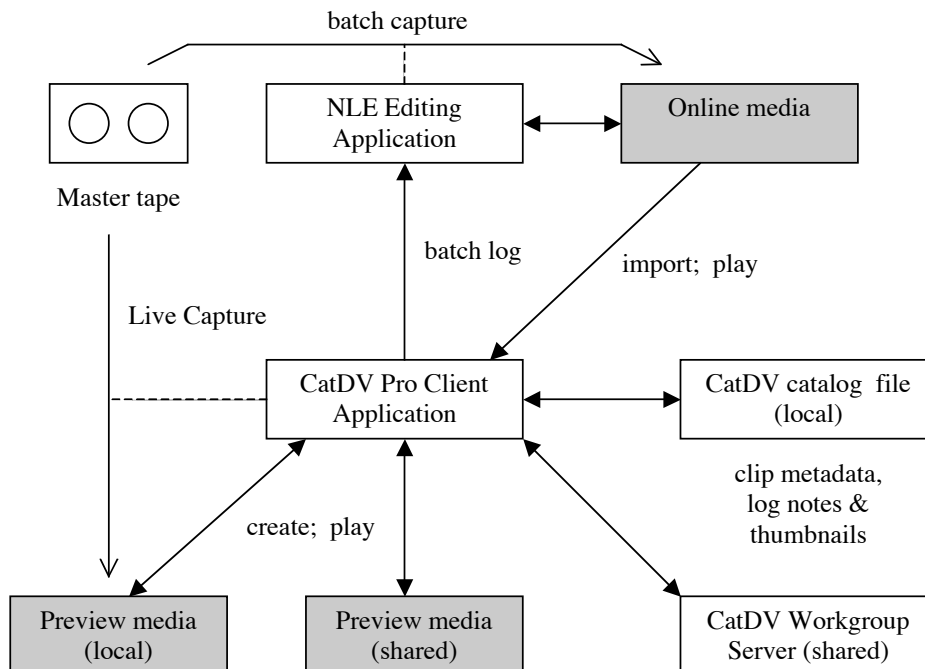
With the *Live HTML Publisher*, the same clip data is further made available through a web interface, allowing it to be viewed from any web browser without requiring the CatDV client software to be installed on that machine. The database can be browsed or searched but in the current implementation this interface is read-only – unlike with the full client application it is not possible to input data via the web interface.

Managing media

The CatDV Pro Workgroup Server does not itself store any media, either full quality on-line media or low-resolution preview versions. Instead, the database primarily deals with metadata *about* the media, such as clip and tape name, video and audio format, log notes, thumbnails, and media file location. This metadata allows clips to be searched and identified, even if the media is currently unavailable.

Once the desired clips have been located using this metadata, the client application can then request the corresponding media. If the media is available online on the local disk this will be via the filename, and the media can be played within CatDV. Alternatively, the user can export a batch capture log of selected clips so the media can be recaptured from tape at full online quality and used within the NLE.

The CatDV Pro Client Application can also create and manage low-resolution *preview* versions of original media. These previews act as “video thumbnails” to give a rough indication of the contents of a clip for those situations when the original full-resolution media are not available online. Although primarily intended for previewing clips within CatDV Pro, the preview files can also be exported and used as low-resolution video editing proxies if so desired. Previews are referenced using a search algorithm based on the tape name and timecode value of the clip, rather than a fixed file or pathname, which means they don’t need to be explicitly associated with a clip but are available for any clip referencing that tape.



Previews are normally stored locally on the client machine, but they can also be shared via the file system by configuring the preview directory to be on a remote volume. However, even though the preview movies are by their nature compressed and much smaller than the online media files, careful thought still needs to be given to both fileservers and network performance if the preview media are to be stored on a remote volume.

If previews are stored on a file server then they can also be made available to web clients via a web server. Again, careful consideration needs to be given to the performance requirements. CatDV does not aim to be a media server by itself but will support this configuration if used in conjunction with a suitable media file server.

Interoperability

As far as possible, CatDV Pro is based around non-proprietary, open standards. The client application is written in Java, and uses Apple’s QuickTime technology, ensuring that it’s available for both Macintosh and PC. The workgroup server and HTML publisher also use Java, and as a result are available for a wide range of server platforms.

By default, the workgroup server uses MySQL as its database. MySQL is a widely available open-source database, available for platforms including Unix, Mac OS X, Windows NT/2000, and Solaris (and included as standard with many Linux distributions). The CatDV Workgroup Server software communicates with the MySQL database using JDBC drivers. By using this open standard it is possible (with some very minor code changes) to install the CatDV database on other vendors' database products rather than MySQL. JDBC also permits the database to be run on the same or on a different machine as the workgroup server.

Similarly, the HTML publisher uses Java servlets, and normally runs within the open-source Apache Tomcat servlet engine, which again is freely available for Windows, Mac OS X and other Unix platforms.

Preview media can be integrated into the web pages returned by the Live HTML Publisher by configuring a web server to return the appropriate files. This might typically be the open-source Apache web server, but other web servers such as Microsoft IIS or dedicated proprietary media servers can be integrated into the system also, as long as the Live HTML publisher is configured to return the appropriate URLs.

To maximise portability and interoperability, CatDV is developed with deliberately conservative assumptions about the platform and libraries it is being run against. Although extensively tested against the latest available versions of third party software (including beta and developer preview versions), care is taken not to rely on later features unless a compelling reason for their inclusion is found. By not depending on the "latest and greatest" (and possibly unstable) versions of software such as Java, QuickTime, MySQL or Tomcat, we can maximise the number of systems on which the CatDV product family will operate.

Toolkit philosophy

The approach taken with the design of the CatDV server software is to use open standards and maintain maximum interoperability and flexibility. While the CatDV Pro client application is a shrink-wrapped end-user application, the server software, especially the HTML publisher, can be viewed as being more of a toolkit, designed to be customised and integrated with other systems as part of the user's overall workflow. It is even possible to license the source code for the CatDV workgroup server itself, permitting CatDV to be hooked in to an existing back end media management system, for example, instead of using its own database.

Workgroup Functionality

Using the workgroup features provided by CatDV Pro several members of a team can share clip catalogs by saving clip information to a central networked database, rather than to files held locally on disk. Users access the central server from the CatDV Pro client application by using commands in the Workgroup menu. This section provides a brief overview of these commands.

Server operation

Before they can use the workgroup server, clients must first *connect to the server* by entering the hostname (and optionally port) of the server machine. The current implementation does not require individual user IDs and passwords (though this is planned for a future update). Instead, each client must be registered with a registration code that is valid for the workgroup server and the total number of connections must not exceed the maximum permitted for that registration code.

Once logged on, the *server status* command will verify the connection to the server and display statistics about the operation of the server, such as how many catalogs and clips are contained in the remote database.

Users can *log off* when they no longer require the server. They are also logged off when the client application quits.

The workgroup server is normally launched via a startup service or from the command line on the server machine. The server can be shut down in the same way, or a *shutdown* can be requested from a connected client.

Publishing and opening catalogs

If the CatDV user has already created catalogs and saved them locally then they need to be *published* before they are available to other users via the shared database. The contents of the current window can be published as a catalog, or an entire folder of catalog files can be published in one operation. In either case, once stored on the server it is no longer necessary to keep a local copy of the catalog file (except as a backup or to support offline working).

Although all the clips on the remote server are stored in the same database, for convenience they are still grouped into logical groupings called *catalogs*. Users should normally create a separate catalog for each tape, or perhaps each shoot or each project, rather than storing everything in one large catalog, as this will make it easier to manage the clips (for example, using the command to *delete* a particular catalog in the database).

The user can *browse* the remote database, displaying a list of all the catalogs including a short summary of their contents, and can then *open* a chosen catalog. Opening a remote catalog creates a local copy in the memory of the client application so the clips can be viewed and manipulated there. Opening a catalog does not lock it, so a user can choose to *refresh* the contents of a window from the database in case another user on the network has made changes in the meantime.

Searching the database

There are two types of search one can perform on the remote database, a *catalog search* or a *clip query*.

Instead of browsing all the catalogs in the database, users can search for catalogs containing a particular keyword. This will search in both the catalog descriptions and the individual clip details. All the catalogs that contain matches are shown and the user may then choose to open one of the matching catalogs in the normal way.

Users can also perform a clip query that searches for clips across the whole database. A window is displayed containing the query results, combining all the clips that match, even if they are in different catalogs. Users can then work with the query results locally, further grouping and filtering the results, printing them out, exporting them as a batch list, or editing the clips and saving the changes back to the database.

Publishing changes and resolving conflicts

When a user opens a remote catalog, or performs a query and is working with the query results, they can edit the clips in their window exactly as if they were working on a normal local catalog file. However, rather than saving these changes to disk, users will normally update the clips in the remote database by using the *publish changes* command.

Users can add logging notes, change clip names, make selections, select new poster thumbnails, delete unwanted clips, split a clip into two or create new secondary clips, and all these changes will be saved back to the database.

If two users try to make changes to the same catalog or clips at the same time then only the first set of changes that are published will be saved. The second person who attempts to publish their changes will receive a warning message stating there were conflicting edits (eg. trying to add a comment to a clip which the previous user has just deleted). The changes aren't lost however.

All the changes which can be saved without conflict are saved, and the main window is refreshed to show the latest contents stored in the remote database. Any clips which weren't able to be saved are displayed in a new "unsaved changes" window. The second user then needs to manually re-apply those changes in the main window, deciding whether and how to resolve any conflicts before trying to publish the changes again.

As well as publishing changes, which merges any changes in with what is already there in the database, it is possible to *overwrite* an entire catalog in the database with a new copy, eg. after working offline on a local copy of the catalog.

Live HTML Publisher

This section describes the functionality provided by the Live HTML Publisher. Note that the HTML Publisher currently only provides a subset of the full workgroup functionality available with the CatDV Pro client application. In particular, it is not possible to import media files or tapes to create new clip records, or to edit existing records, using the web interface.

Overview

The Live HTML Publisher consists of Java servlets and JSPs (Java Server Pages) that communicate with the CatDV MySQL database to display the contents of the database. The interface is based around separate web pages. The following pages are available as standard:

- Browse the database, displaying a list of all the catalogs in the database together with a brief summary of the contents of each catalog (dates, tapes, and clip types included in that catalog).
- Provide details of a particular catalog, including a list of all the clips in that catalog.
- Provide details of a particular clip, including thumbnail, a list of all the available clip properties, and a link to the preview media if available.
- Enter parameters for a search based on text fields or dates and display the results of a search, either as a list of matching catalogs or of matching clips.
- Display statistics about the database (total number of clips stored, etc.)

Customising the interface

Java Server Pages are HTML pages that contain special tags that execute Java code on the server to query the CatDV database and insert data into the page. (JSPs are processed by the Apache Tomcat servlet engine. JSPs are automatically compiled into Java servlets when they're accessed for the first time). It is possible therefore to use normal HTML editing tools to customise the layout of JSP pages and add new functionality (though it is important to get the syntax of the tags correct otherwise you will get Java errors when you display the page!)

Although functional as provided, the pages that come with the Live HTML Publisher are intended as templates that can be customised by the user as required.