

HardingFPA



CAMBRIDGE RESEARCH SYSTEMS

HardingFPA-X Administration Manual

This manual details how to install and set up the HardingFPA-X system.

For HardingFPA-X Version 3.2.0

Documentation date: 12/12/2012.

Overview

The HardingFPA-X System is a distributed version of the HardingFPA Flash and Pattern Analyser. It operates in a similar way, but is intended for file-based workflows instead of tapes.

If analysing HD material, the HardingFPA-X analyses using the new Version 3 analysis algorithms, which have been designed specifically for file analysis in the HD world. For compatibility with older algorithms the software can be set up to use SD Legacy Mode, which allows SD material to be analysed using Version 2.5 analysis algorithms which are the same as in the previous 2.54/2.57 versions of the HardingFPA, and Version 1.x of the HardingFPA-X. HD material is always analysed using Version 3 analysis algorithms.

This manual gives an overview of the components in the HardingFPA-X System, and details the installation and configuration process for the system. It does not include detailed user instructions for the Client side applications (HardingFPA-X *Submit* and HardingFPA *Viewer*). These can be found in the HardingFPA-X User's Manual.

Jobs are submitted either via client side interfaces (*Submit* applications) or by dropping the movie files into watch folders. These jobs are then posted to the central database into the Job Queue, where they are picked up by analysis installations (running the *Analyser* application) which analyse the material, typically at a greater speed than real-time, and write out results which the client users may then view. The *Analysers* will analyse jobs of a greater priority first, and analyse the earliest submitted jobs first. Only the system administrator has the ability (via the *Supervisor* application) to modify the priorities of jobs and abort jobs.

The HardingFPA-X System is provided as three separate installer distributions:

- **Server**, which contains the components required to run the analysis and watch folder services.
- **Supervisor**, which is an administrator's tool used for licensing clients, viewing system status and managing the Job Queue.
- **Client**, which can be installed on multiple licensed computers, allows users to submit clips to be analysed and view results across a local network connection.

It is entirely possible to install all of the distributions on the same computer (Server, Supervisor and Client), but better performance will be gained if the Server components are run on a dedicated computer by themselves, preferably a high-performance multi-core computer.

Optionally, a Web Services interface can be provided via SOAP, using modules for Apache or ISS web-servers (please see the HardingFPA-X Programmer's Manual for more details).

The HardingFPA-X components are available for both Windows and OS X. Within your HardingFPA-X System you may “mix and match” computers running the OS X or Windows versions of the components, in whatever combination you wish.

This manual covers both the OS X and Windows versions of the HardingFPA-X software.

The versions of the components in the current release are detailed in the table below. Please refer to the *Changes.html* file that is provided with the HardingFPA-X install packages for details on changes and new features present in this release.

Component	Version	Purpose	Distribution
HfpaPostgreSQL	3.2.0	The main database component, which all other components in the HardingFPA-X System connect to. This database maintains the job lists, licensed clients and results locations. Only one of these is required per HardingFPA-X System.	Server (OS X) SQL Script provided for Windows
Launcher	3.2.0	Maintains and ensures that the Analyser module continues to run normally at all times.	Server
Analyser	3.2.0	Analyses movie clips submitted to it via users or from watch folders.	Server
Monitor	3.2.0	Manages watch folders. Continually scans single or multiple watch folders, submitting movies for analysis. Also facilitates advanced XML-based job submission.	Server
Supervisor	3.2.0	Allows licensing clients, viewing of system status and administrator manipulation of the job queue via the use of job priorities as well as the option to abort jobs.	Supervisor
Submit	3.2.0	Provides the interface to the job database for client users. Allows submission of new jobs and viewing of results of completed jobs.	Client
Viewer	3.11.0	Shows graphical and numerical representations of the results of completed jobs to users on a frame-by-frame basis.	Client

Please note that detailed information regarding the operation of the Client components can be found in the HardingFPA-X User's Manual.

System Requirements

The HardingFPA-X System is cross platform, and each component can run on both Windows and Mac computers.

The applications will only run on Intel powered Mac computers, and require OS X version 10.5 or higher for all components. Computers running the Analyser component will obtain optimum performance with multi-core processors, with at least 2 GB of RAM.

For Windows platforms, the HardingFPA-X has been tested on Windows XP, Windows 7, Windows Server 2003, Windows Server 2008. Although the HardingFPA-X is compatible with Windows Vista, this operating system is not recommended. At least 2 Gb of RAM is recommended for the best performance for the computer running the Analyser.

Hint: It is only possible to install the PostgreSQL (database) component of the HardingFPA-X System using the automated Mac installer on a Mac computer which does not already have an installation of PostgreSQL (except for upgrades to the HardingFPA-X PostgreSQL). If you have another application which installs a PostgreSQL database, for example, you may need to install the HardingFPA-X Database on a different computer or manually install the roles and tables for the HardingFPA-X System from the provided hfpa.sql file.

Hint: The PostgreSQL Database may be limited in the number of connections it may be able to accept, which could impact on the usability of the HardingFPA-X System. If this is the case, then you may install the Database component on a separate computer running a Server version of the operating system (OS X Server or Windows Server 200x) or import the hfpa.sql file into a PostgreSQL installation running under a distribution of Linux.

Licensing

Computers running the *Analyser* must have a HASP USB software protection key (see below), and will not operate without one.

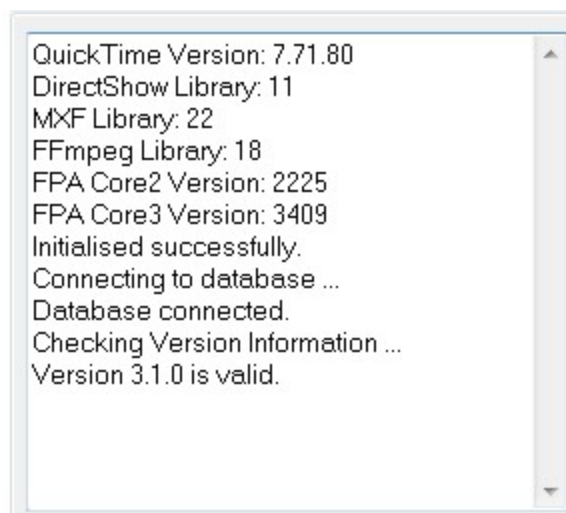


Computers running the *Submit* application will need to be registered in the database using the HardingFPA Administrator password (See [Licensing Clients](#)) before they are allowed to run.

Version Control

The HardingFPA-X System features an automatic version control system to help you keep your system fully up to date. When components start up, they will check with the database to see if their version numbers are valid (i.e. not older than the latest version currently in use).

If the version is not valid the applications will advise the user of this with a message dialogue box and then shut down. In this case you should upgrade all the components in your system to the same version as the newest version.



If an application finds, whilst starting up, that it is newer than the rest of the system, the version for the system is updated to the new version number, and the older applications will cease to work until they too are updated.

Movie Files

The system will open and attempt to analyse any movie file for which the computer it is installed on has the codec. There are number of codecs included (see below). Any further codecs required will need to be installed on the computer(s) running the application.

Only movie files with the following extensions will be analysed:

***.avi, *.mov, *.mpg, *.mpeg, *.m2v, *.mp4, *.vob, *.wmv, *.mxf, *.flv, *.qt, *.ps, *.3gp, *.mkv, *.m2ts**

In general, if the file can be viewed correctly using QuickTime (or Windows Media Player) on the computer that the application is running on, and is within the accepted limits for frame size and frame rate (see below), then the application will be able to analyse it.

The list of accepted frame sizes and frame rates are as follows:

Legacy SD analysis frame sizes:

320x240, 352x288, 384x288, 640x480, 640x486,

702 – 720 x 480

702 – 720 x 486

702 – 720 x 576

768x576

Legacy SD analysis frame rates:

25, 29.97

Version 3 SD analysis frame sizes supported:

320x240, 352x288, 352x240, 384x288, 480x360, 640x360, 960x540 640x480, 854x480, 640x486,
854x486,

702 – 720 x 480

702 – 720 x 486

702 – 720 x 576

768x576, 1024x576,

Version 3 HD analysis frame sizes supported:

960x720, 1280x720, 1440x1080, 1920x1080

Version 3 analysis frame rates:

24, 25, 29.97, 30, 50, 59.94, 60

The HardingFPA-X uses the following frameworks to access video frames:

- FFmpeg
- QuickTime
- DirectShow (on Windows only)

When presented with a file, a framework is selected based on the movie file extension.

MXF Files

An attempt is made to open the MXF file using the internal MXF library. This library has support for OP1A and OP ATOM wrapped MXF files with the following codecs:

- D10 (IMX)
- DV
- DVCPRO
- DVCPRO HD

- DNxHD
- AVC Intra

If the internal MXF library does support the MXF file format, then the FFmpeg library is used to decode these frames. The MXF library includes support for AVID MXF files as well as MXF time code tracks.

If the MXF file format is **not** supported by the internal MXF library, an attempt is made to use the FFmpeg library to open the file directly.

On Windows platforms, if all else fails an attempt to use DirectShow (using any codecs installed in the system) will be made. If this is successful, an attempt will be made to use the VITC information, if any, embedded in the video material.

MOV Files

The QuickTime library will be used to open QuickTime files. File support depends on the QuickTime codecs that are installed on the computer used for analysis. The QuickTime library supports QuickTime timecode tracks. The system does not support multi-track QuickTime movies.

Other Files

Firstly the use of FFmpeg is attempted, followed by an attempt (on Windows platforms only) to use the DirectShow framework, and finally the QuickTime framework.


Why use FFmpeg?

- FFmpeg provides a uniform set of embedded codecs that allow for consistent results across different platforms and installations.
- If a previous version of the application used DirectShow or QuickTime, and a newer version uses FFmpeg, there is a possibility that the results will differ due to the slight differences between codec algorithms. The use of FFmpeg in current and future versions of the HardingFPA product range will mitigate this problem.
- FFmpeg in many cases removes the requirement of purchasing 3rd party MXF support and codecs as the MXF library, in conjunction with FFmpeg, now fulfils this requirement.

The FFmpeg library supports the following codecs:

D10 (IMX), DV, DVCPRO, DVCPRO HD, DNxHD, MPEG1VIDEO, MPEG2VIDEO, H261, H263, RV10, RV20, MJPEG, MJPEGB, MPEG4, MSMPEG4V1, MSMPEG4V2, MSMPEG4V3, WMV1, WMV2, H263P, H263I, FLV1, H264, INDEO3, VP3, THEORA, ASV1, ASV2, FFV1, MSVIDEO1, SNOW, XVID, FFVHUFF, INDEO2, FRAPS, VP5, VP6, VP6F, FFH264.

The decision tree that is used to decide which framework to use is shown in the flowchart below:




Installation (OS X)

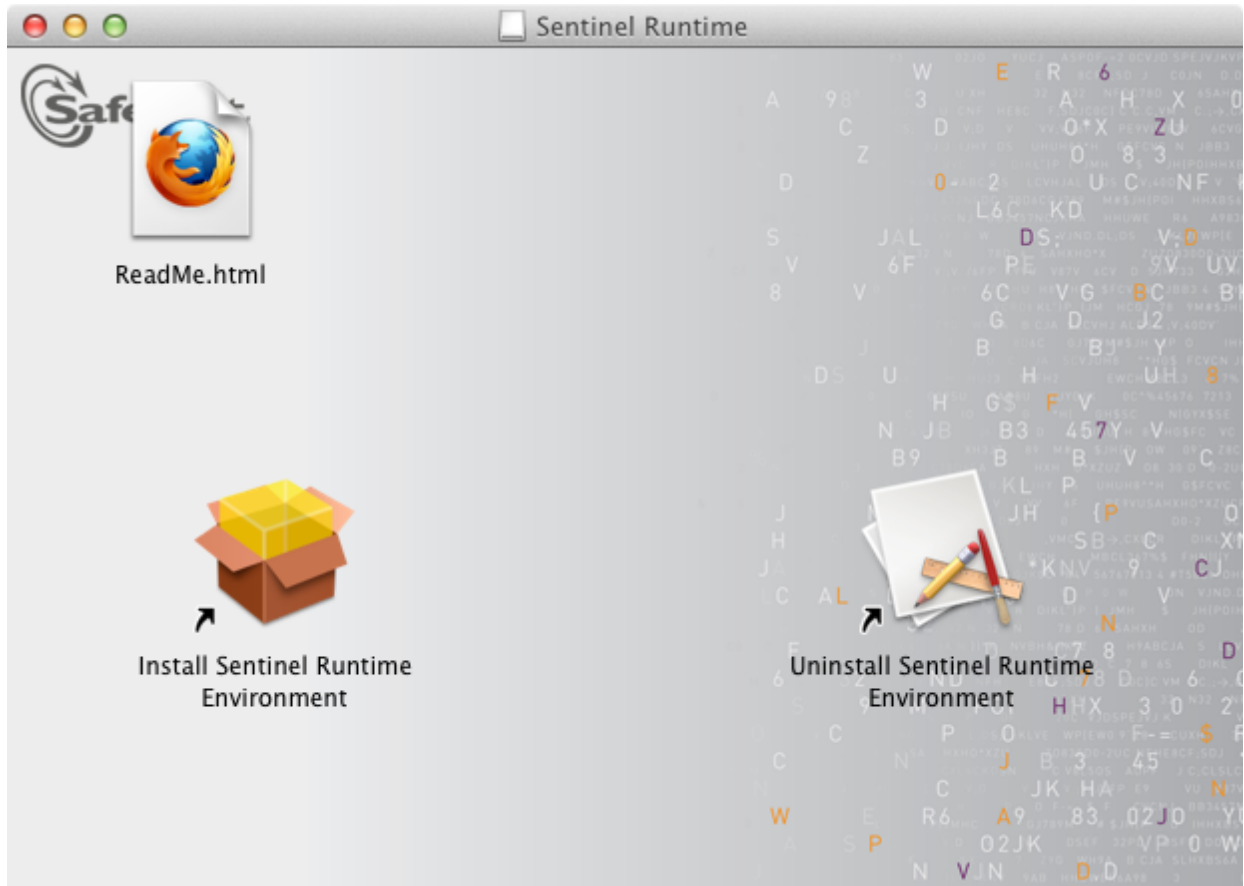
This section describes the installation of the HardingFPA-X components on the OS X platform.

Prerequisites (OS X)

In order for the *Analyser* parts of the Server distribution to work correctly, a USB HASP key must be installed, and therefore the HASP drivers must be installed. To install these, double-click on the *Sentinel_Runtime.dmg* disk image (shown below), which can be found inside the HardingFPA-X disk image.




This will mount the disk image containing the driver installer. This screen can be seen below. Double-click on the *Install Sentinel Runtime Environment* icon within this image to install the driver.



Server Installation (OS X)

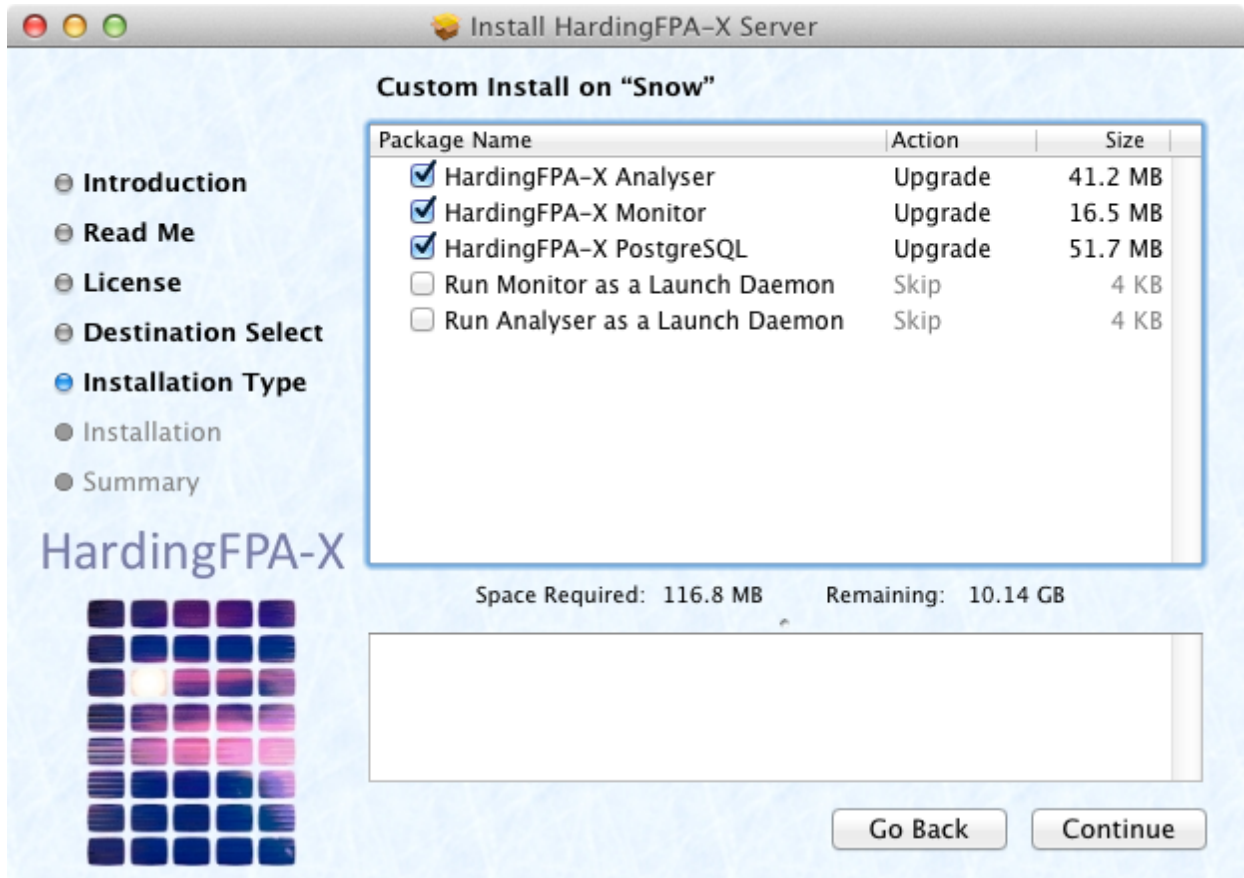
- Double-click on the *HfpaServer* installer package as shown below:



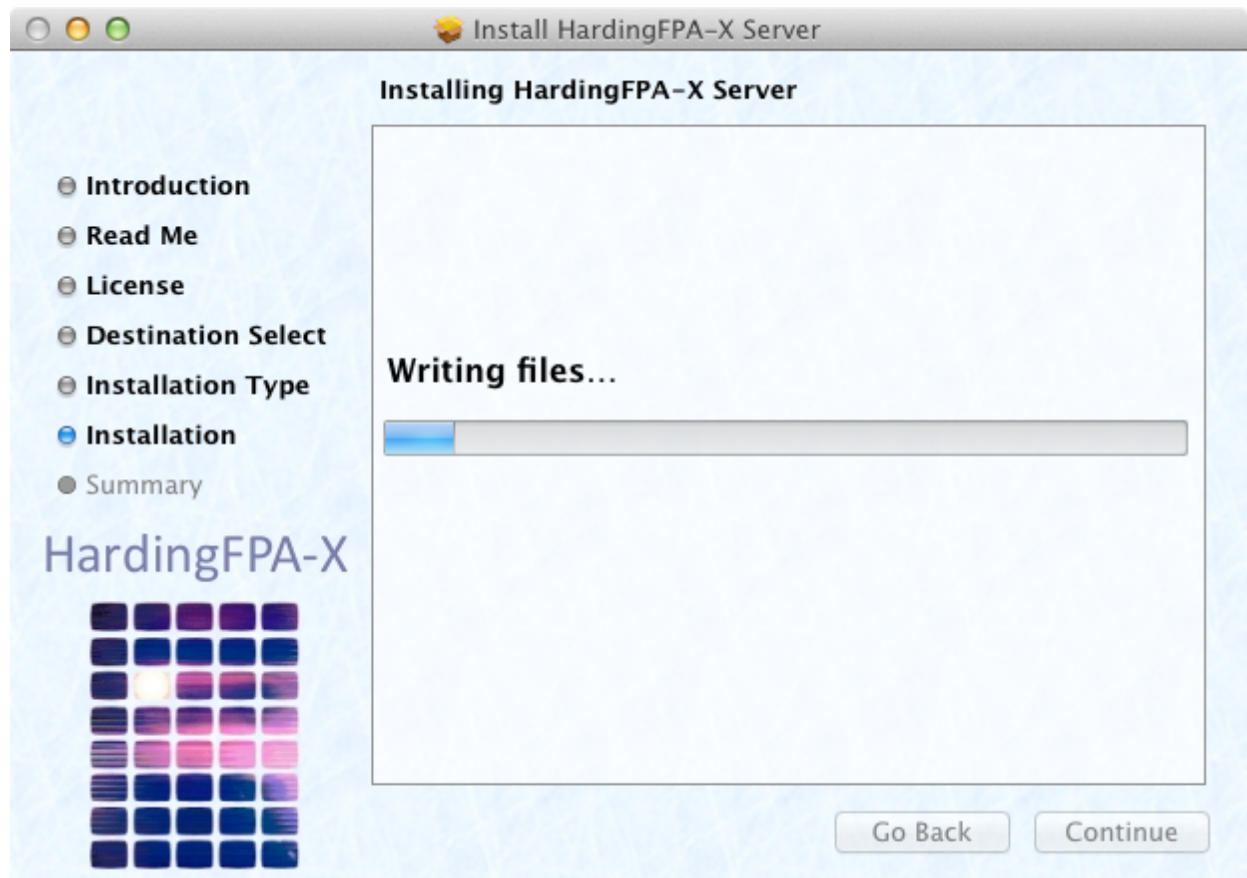
- The installer will then start and the following dialogue box will open:




- Click on continue, then read the Read Me text and accept the licence terms displayed on the next screen.
- Select which volume (hard disk) you would like to install the server applications to.
- Select which of the components you would like to install. The *PostgreSQL* component is a database that is required for the HardingFPA-X system to run correctly, but will typically be installed on only one of the computers on your network, so if you are installing a second or third *Analyser* node for example, you should de-select this option.
- To run the *Monitor* or *Analyser* without logging in, select the option(s) to run as a Launch Daemon. Please refer to the separate user manual on running HardingFPA-X applications as Launch Daemons for more information on how to manage this.



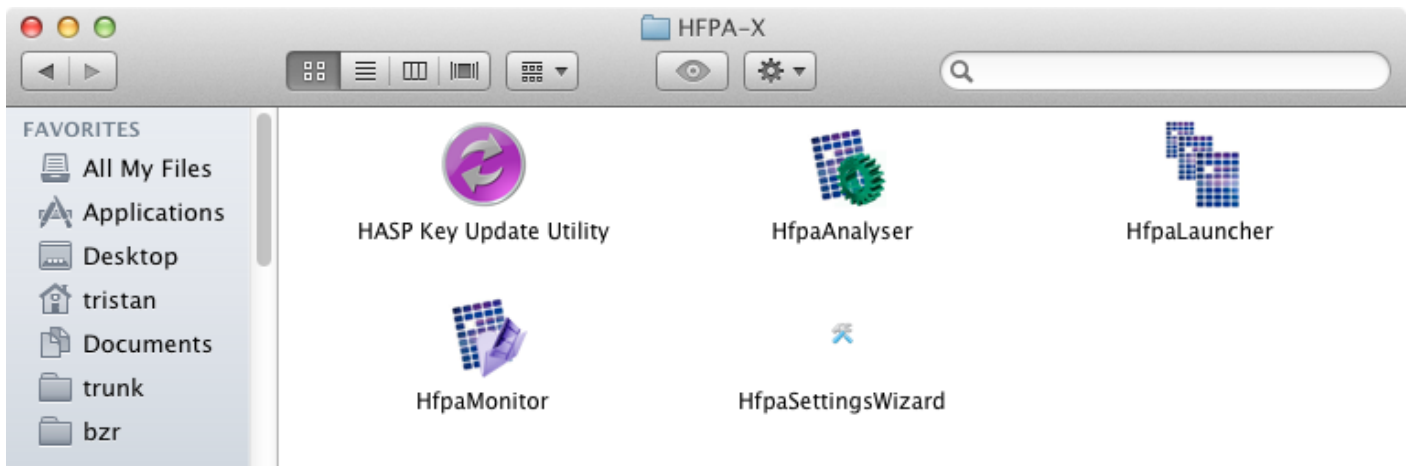
- Wait for the installation to complete.



- After completion, the PostgreSQL (database) configuration utility will appear. If you only require the system to be run from the same computer, and have no remote clients, then choose Local Host. If you would like to be able to connect client computers from across your network, then select the appropriate IP addresses for your local network cards by checking the relevant boxes.




- It is usually sufficient to leave the port number set to 5432, but if you know that this port is being used by another service or application, it is possible to customise the port that the database runs on by entering a new port number into the box provided. Please make a note of these settings in order to configure the other components in the system.
- Click OK to complete the installation.
- The installed applications will appear under Applications → HFPA-X in Finder, as shown below:



Client Installation (OS X)

- Double-click on the *HfpaClient* installer package shown below:




- The installer will then start.
- Select which volume (hard disk) you would like to install the client applications to.
- Click on continue and wait for the installation to complete. There is no configuration utility for the client side software during the install process, so once this has completed, all the client components will be installed under Applications → HFPA-X in Finder.

Supervisor Installation (OS X)

Note: If the HardingFPA-X Supervisor is installed on a the same computer as the database, then the password login is disabled, and any user running it will be able to login and therefore licence clients and change the priorities of jobs in the system and abort jobs already running.

- Double-click on the *HfpaSupervisor* installer package shown below:




- The installer will then start.
- Select which volume (hard disk) you would like to install the HardingFPA-X Supervisor.
- Click on continue and wait for the installation to complete. There is no configuration utility for the Supervisor software during the install process, so once this has completed the HardingFPA-X Supervisor will be installed under Applications → HFPA-X in Finder.


Installation (Windows)

This section describes the installation of the HardingFPA-X components on the Windows platform.

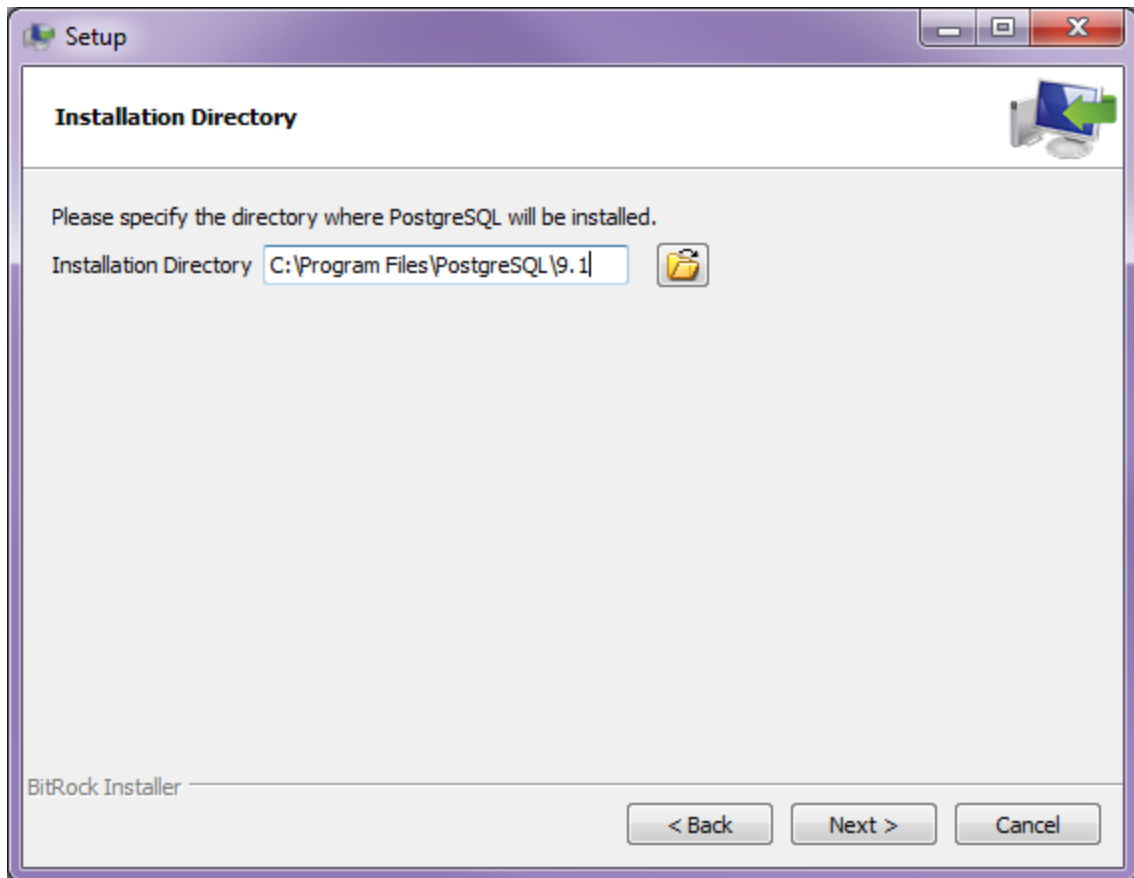
Database Installation (Windows)



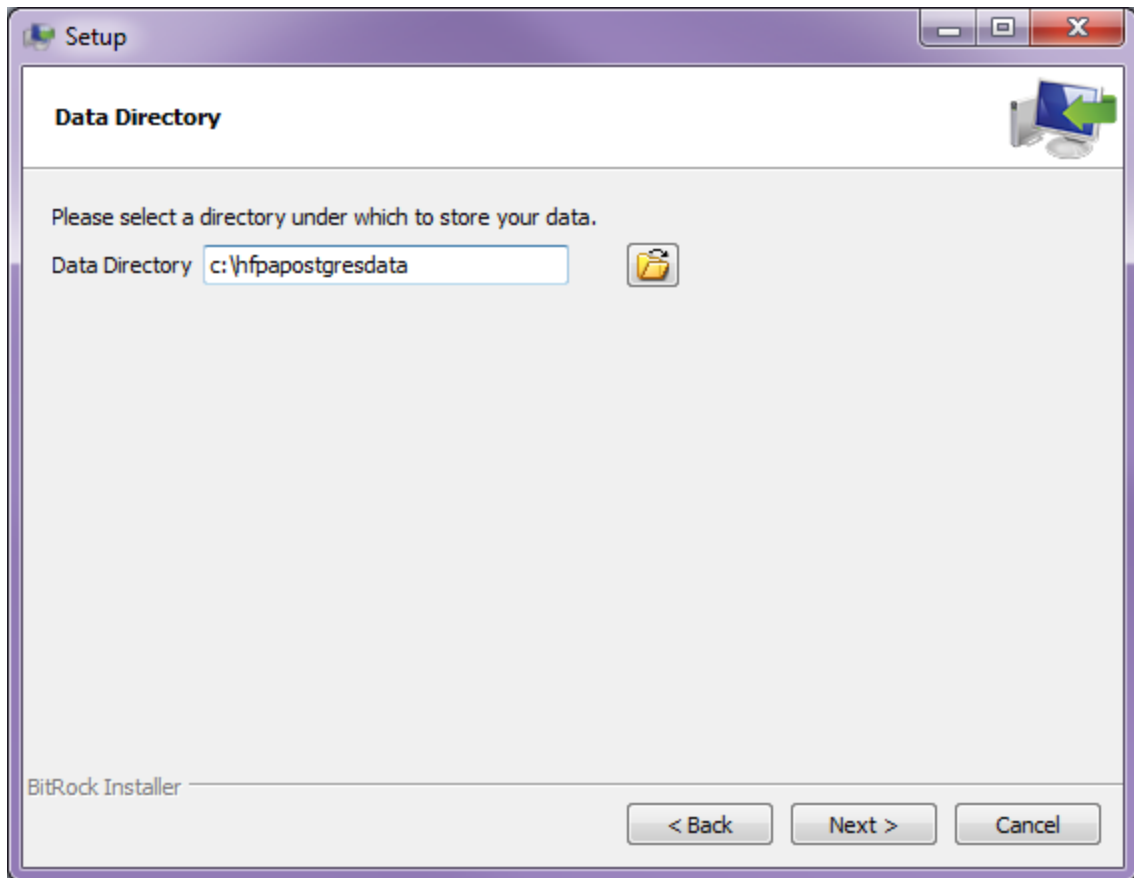
- Double-click the *postgresql-9.1.4-1-windows.exe* file (shown above), which can be found on the installation media. The following screen appears, click "Next":



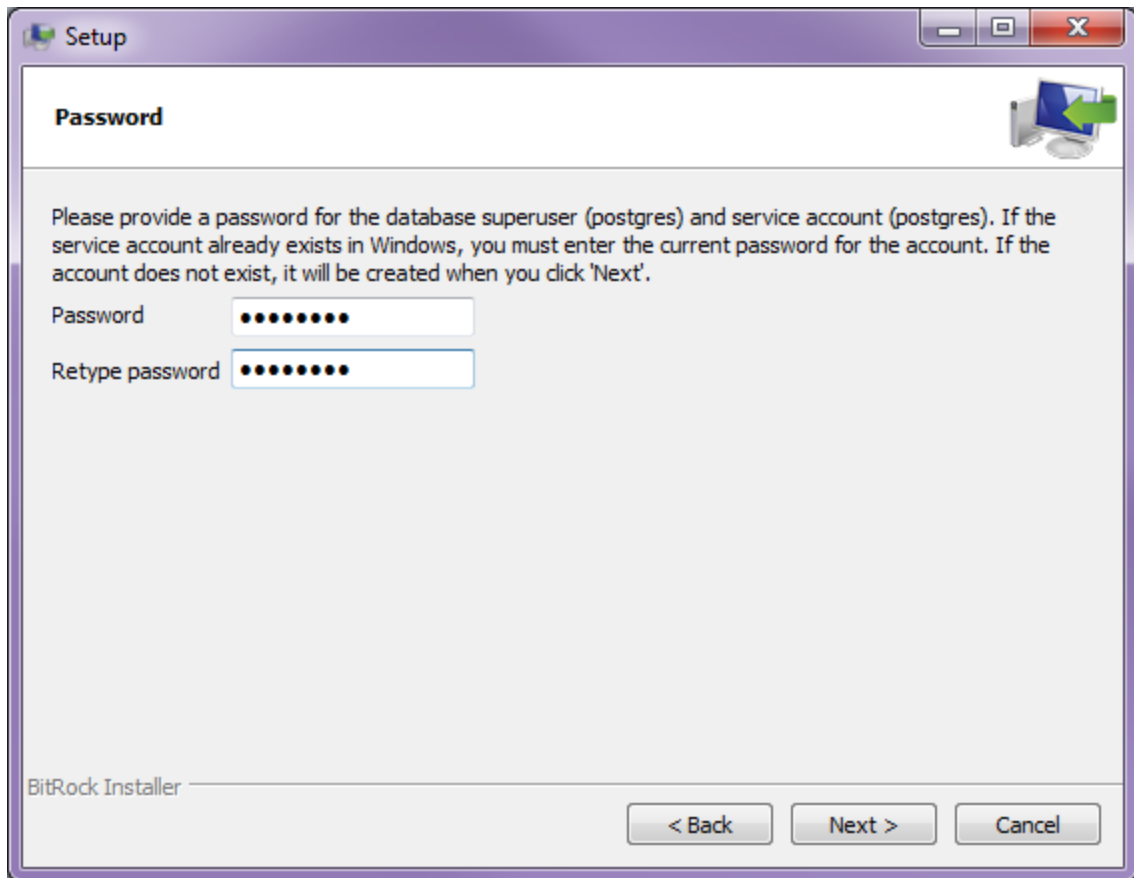
- On the next screen, choose the installation directory and click "Next" (this can be left at the default value):




- On the next screen, change the data directory to a location on your main drive (usually C), rather than a location within Program Files, to avoid permissions problems. The recommended location is *c:\hfpapostgresdata*



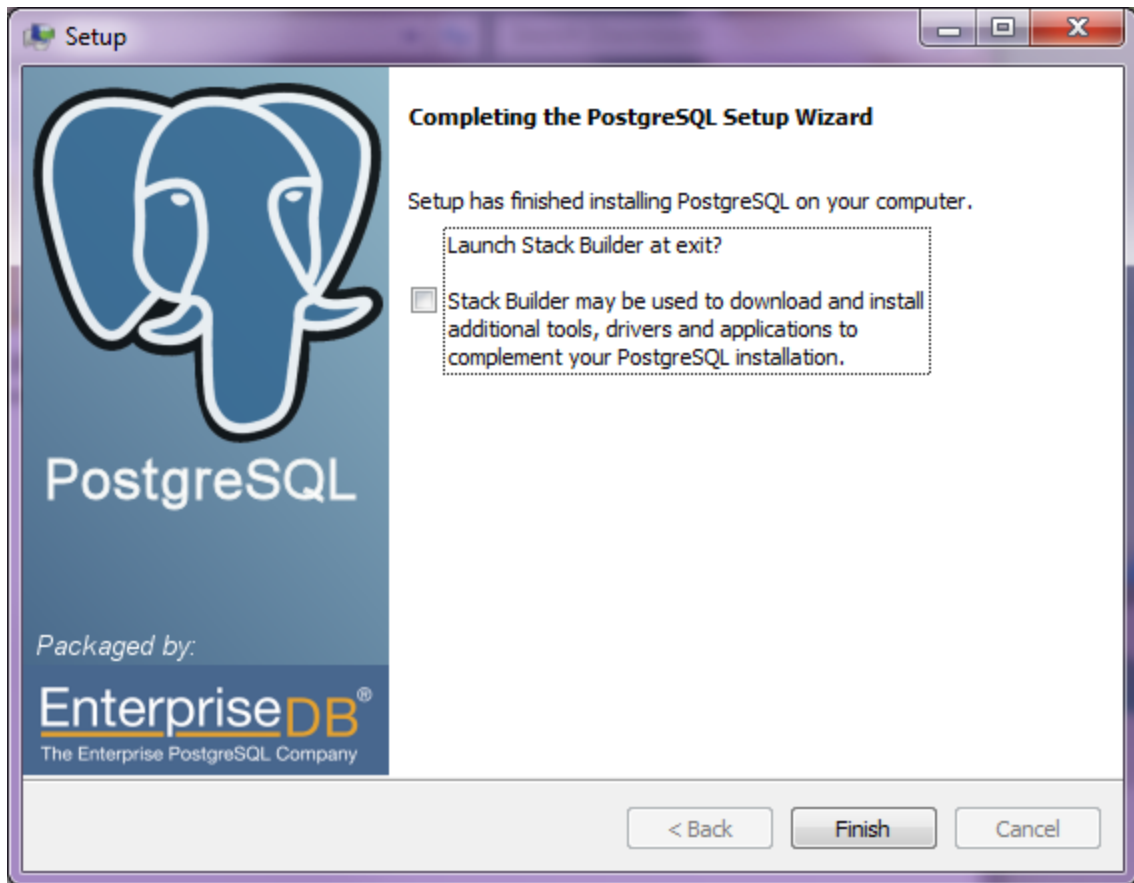
- On the password screen, set a password for the new PostgreSQL Windows user, as shown below. Please remember this password as it may be used in future database upgrades. If upgrading the hfp database, please use the same password as last time:




- Click "Next", then "Next" again to start installing. See below:



- The final screen (shown below) will appear. Uncheck the box "Launch Stack Builder at Exit" and click "Finish".



- Copy the *hfpa.sql* file from the HardingFPA-X installation media to the C:\temp folder of your computer.
- Under *PostgreSQL 9.1* on the start menu, open (click) the SQL Shell (psql) menu item. This will open the command prompt shown below:



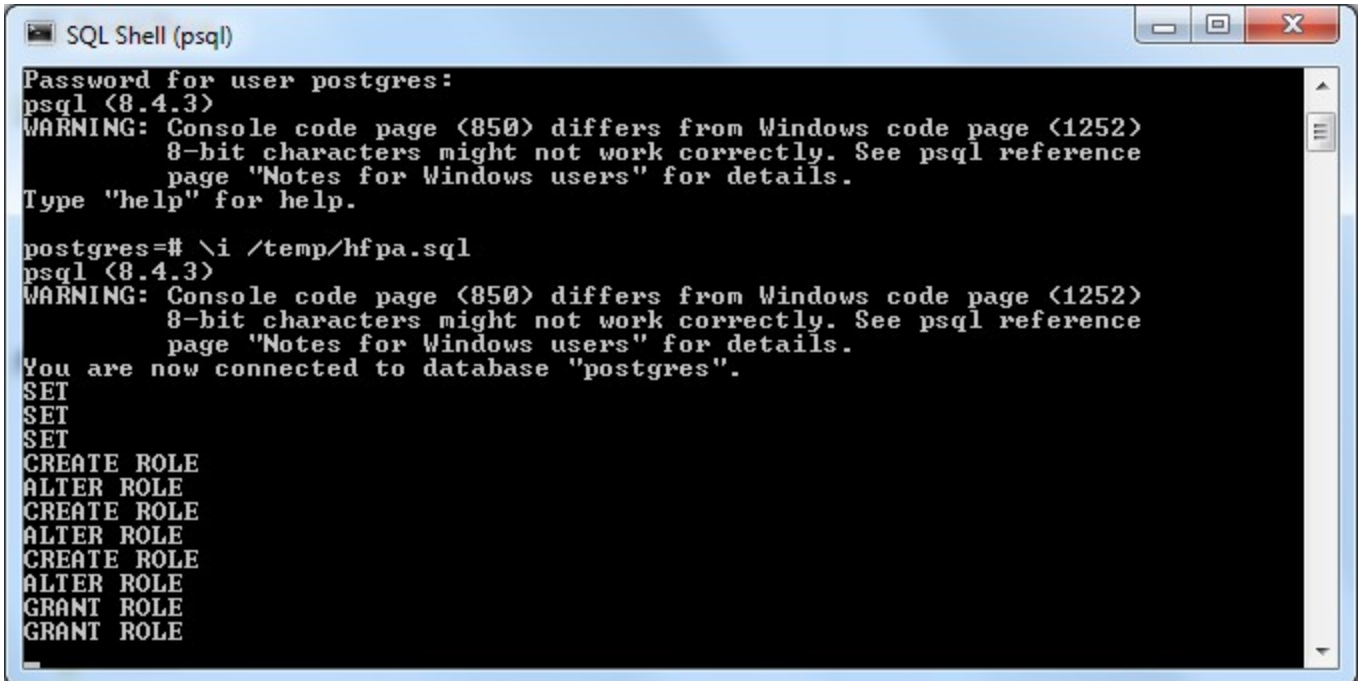
- Press Return 4 times.
- When at the screen below, enter the password you set earlier for the postgres Windows user to log in to the SQL shell.



- Once logged in, enter the following command:

```
\i /temp/hfpa.sql
```

and press Return. The roles and databases for the HardingFPA-X system will be created (see below):



```
SQL Shell (psql)
Password for user postgres:
psql (8.4.3)
WARNING: Console code page (850) differs from Windows code page (1252)
8-bit characters might not work correctly. See psql reference
page "Notes for Windows users" for details.
Type "help" for help.

postgres=# \i /temp/hfpa.sql
psql (8.4.3)
WARNING: Console code page (850) differs from Windows code page (1252)
8-bit characters might not work correctly. See psql reference
page "Notes for Windows users" for details.
You are now connected to database "postgres".
SET
SET
SET
CREATE ROLE
ALTER ROLE
CREATE ROLE
ALTER ROLE
CREATE ROLE
ALTER ROLE
GRANT ROLE
GRANT ROLE
```

- In order to allow PostgreSQL to accept connections from hosts other than the local host, the file named `pg_hba.conf`, which is located in the data directory you chose whilst installing e.g:

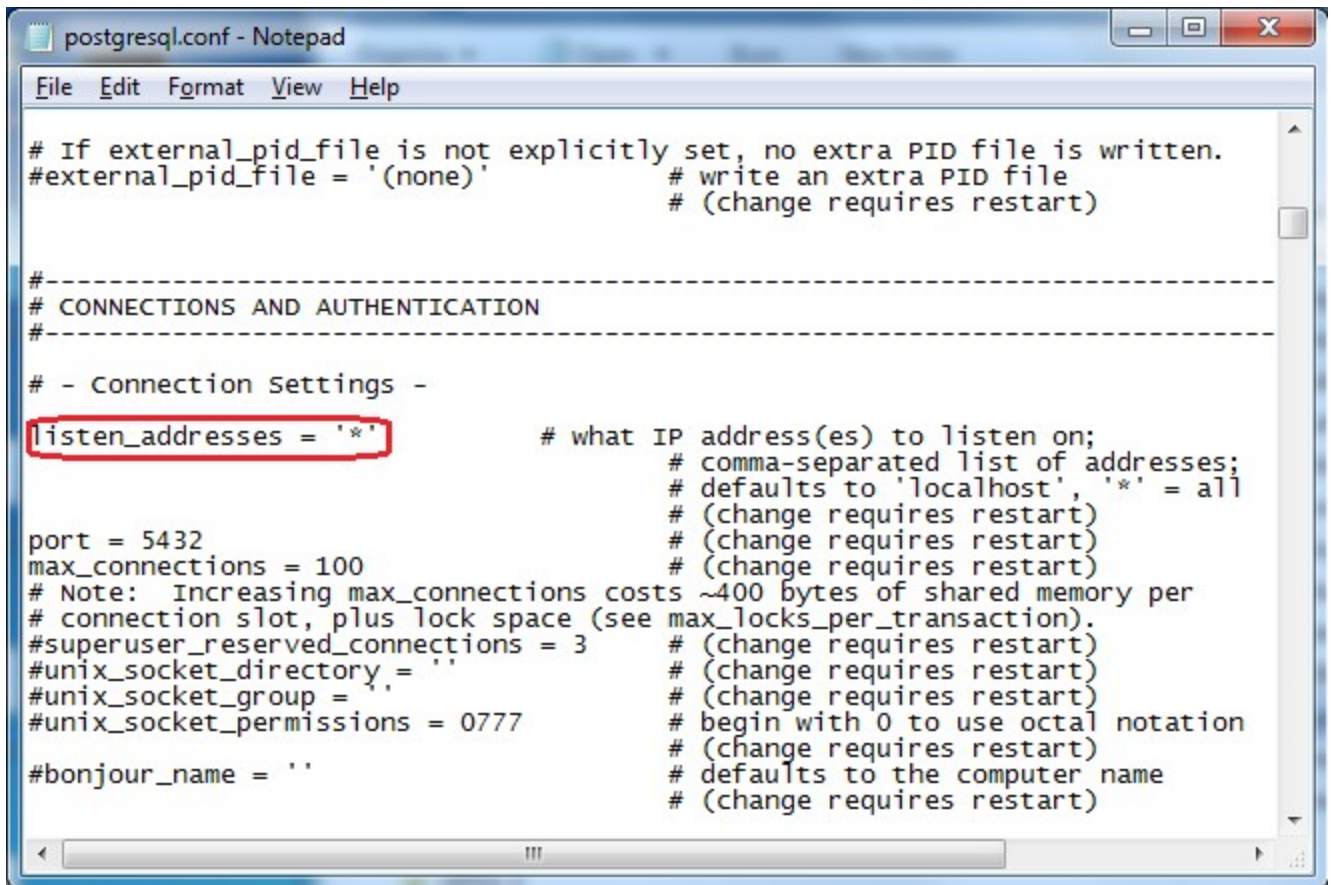
```
c:\hfpapostgresdata\pg_hba.conf
```

must be edited by adding the following line to the bottom of the file.

```
host all all 192.168.0.0/24 md5
```

Please note that the [IP address/subnet mask combination \(in CIDR notation\)](#), for example [192.168.0.0/24](#) in the line above, should match your network's subnet range.

- Make sure that the `listen_addresses` are set to `'*'` in the file `c:\hfpapostgresdata\postgresql.conf`. (or wherever you chose to install the data directory to). See below:



```

postgresql.conf - Notepad
File Edit Format View Help
# If external_pid_file is not explicitly set, no extra PID file is written.
#external_pid_file = '(none)'          # write an extra PID file
                                       # (change requires restart)

-----
# CONNECTIONS AND AUTHENTICATION
-----


# - Connection Settings -
listen_addresses = '*'                # what IP address(es) to listen on;
                                       # comma-separated list of addresses;
                                       # defaults to 'localhost', '*' = all
                                       # (change requires restart)
port = 5432                            # (change requires restart)
max_connections = 100                  # (change requires restart)
# Note: Increasing max_connections costs ~400 bytes of shared memory per
# connection slot, plus lock space (see max_locks_per_transaction).
#superuser_reserved_connections = 3    # (change requires restart)
#unix_socket_directory = ''            # (change requires restart)
#unix_socket_group = ''                # (change requires restart)
#unix_socket_permissions = 0777       # begin with 0 to use octal notation
                                       # (change requires restart)
#bonjour_name = ''                    # defaults to the computer name
                                       # (change requires restart)

```

- Restart the computer.
- You may need to configure your firewall to allow for connections to the PostgreSQL database from other computers, otherwise HardingFPA-X applications running on other computers may not be able to operate.

Server Installation (Windows)

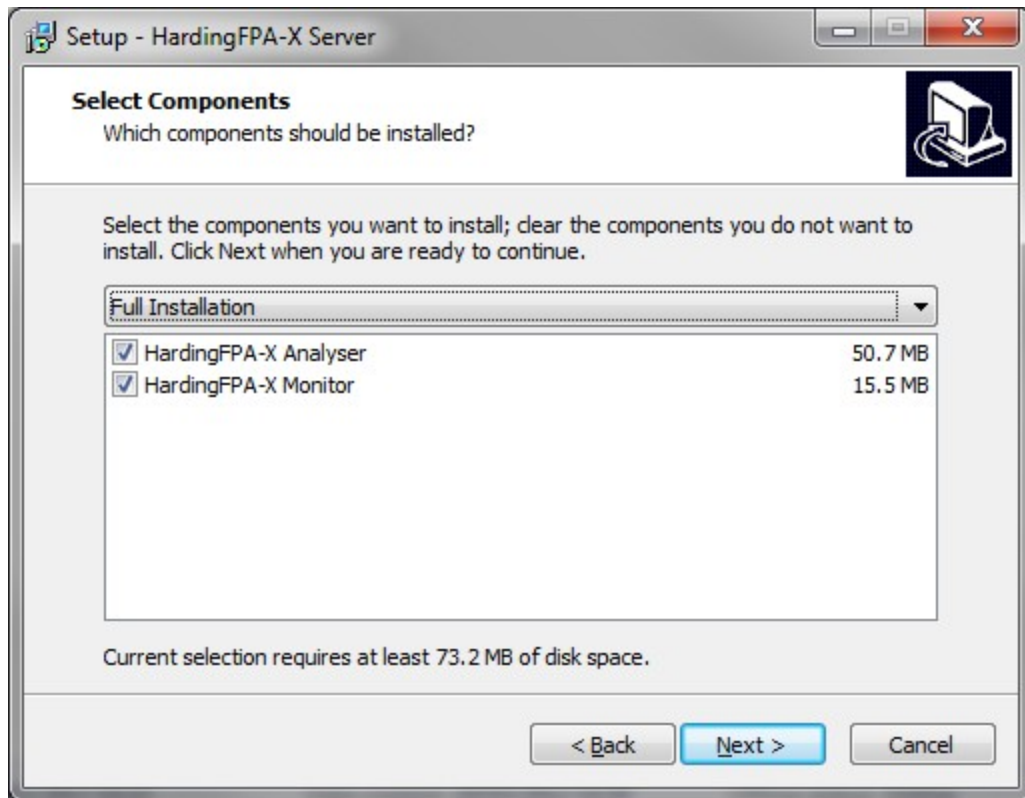
Note that the Windows installation of the Server components do not by default include a database component. Please see the section on [Database Installation \(Windows\)](#) for more information.



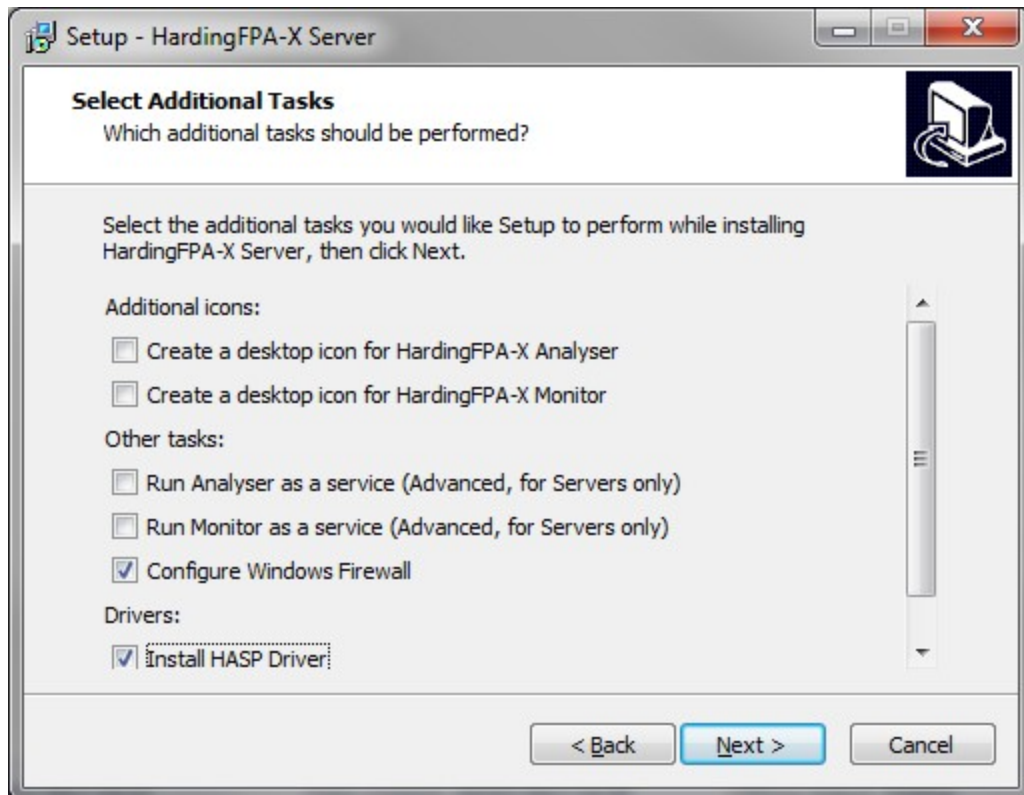
- Double click on the *HardingFPA-X Server Setup.exe* installer shown above. The installer screen below will then be displayed.



- Click Next. Read and accept the licence terms. If you cannot accept these licence terms, please cancel the installation and contact Cambridge Research Systems Ltd.:
- On the next screen select where to install the software. The screen below will appear, allowing you to install only the components you require.




- Click Next, and select whether or not to include shortcuts on the Desktop. *Install HASP Driver* is only checked if it is not already installed previously, but can be checked manually on this page should you wish to reinstall it.
- If you would like to run the *Monitor* and/or *Analyser* as a service (meaning that no interface is visible, and the software can run without a user being logged in) then select the appropriate options now. Please refer to the separate manual on how to operate the system as a service on the installation media.

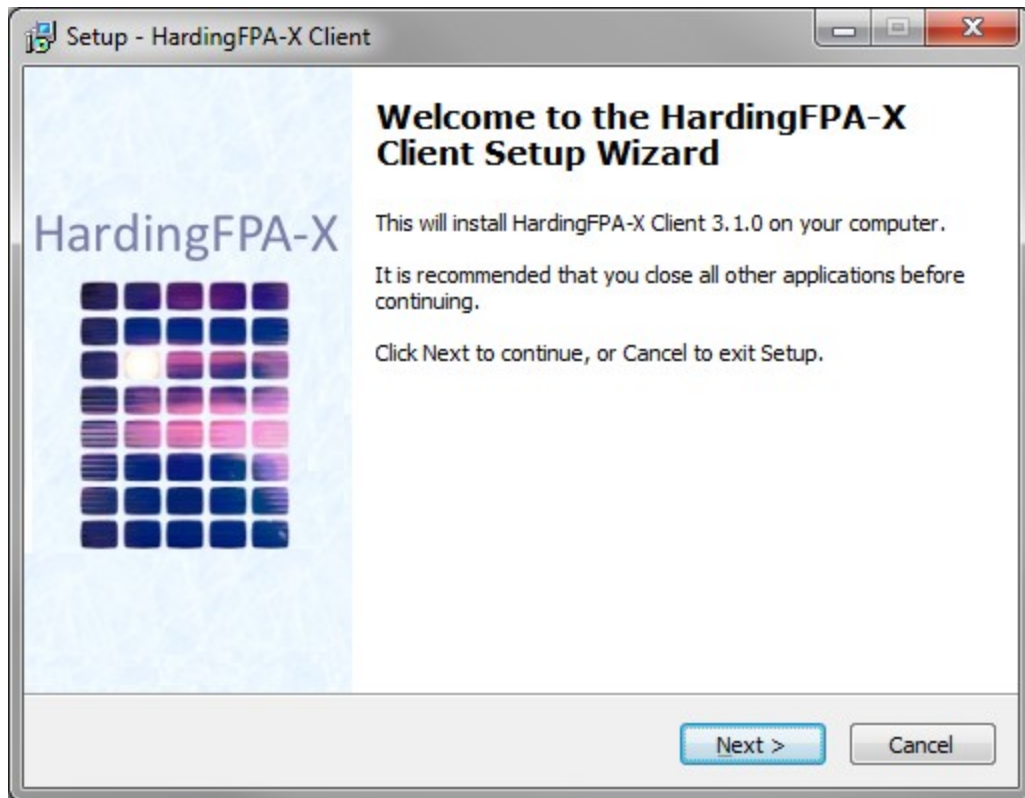


- The installer will then install the chosen applications (see below). The required HASP drivers and runtimes will also be installed at the same time, and may require you to restart the computer on completion.
- After clicking Finish, the chosen applications will have been fully installed, and will appear (by default) under the HSPA-X group on the Start Menu.

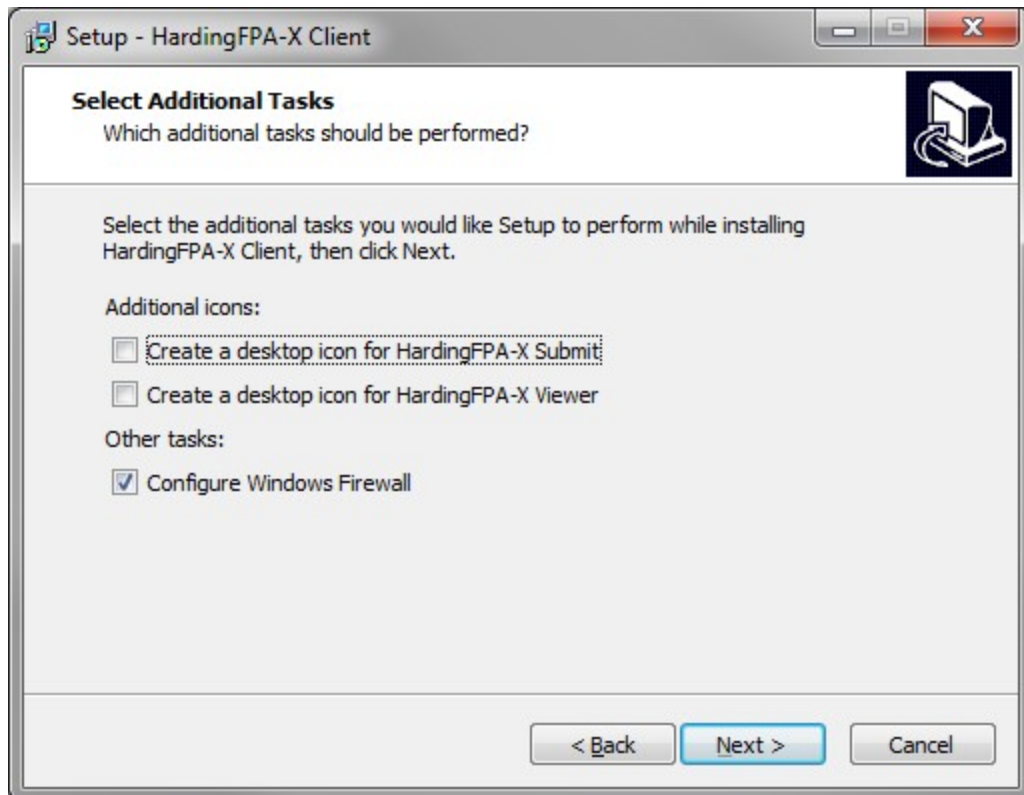
Client Installation (Windows)



- Double-click on the *HardingFPA-X Client Setup.exe* installer package shown above.
- The installer screen below will then be displayed.



- Click Next. Read and accept the licence terms. If you cannot accept these licence terms, please cancel the installation and contact Cambridge Research Systems Ltd.:
- On the next screen select the location into which the software is to be installed, and whether to include desktop shortcuts.
- In order to operate properly, the client applications need to make connections across the network to the database which may be blocked by the Windows Firewall. The applications use numerous ports automatically depending on the computer's configuration. To automatically configure the firewall, click the "Configure Windows Firewall" option shown below.



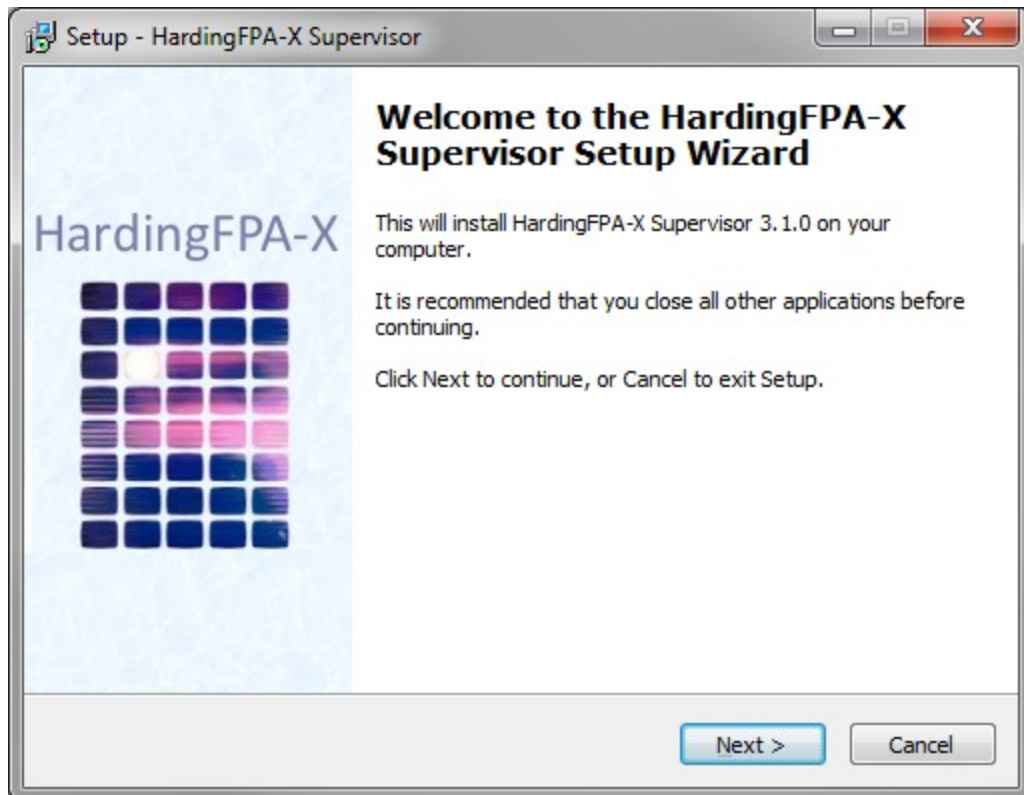
- Click Next, then click Install. Wait for the installation to complete.
- After clicking Finish, the HardingFPA-X *Submit* and HardingFPA *Viewer* will have been fully installed, and will appear (by default) under the HFPA-X group on the Start Menu.

Supervisor Installation (Windows)

Note: If the HardingFPA-X Supervisor is installed on a the same computer as the database, then the password login is disabled, and any user running it will be able to login and therefore licence clients and change the priorities of jobs in the system and abort jobs already running.



- Double-click on the *HardingFPA-X Supervisor Setup.exe* installer package shown above.
- The installer screen below will then be displayed.



- Click Next. Read and accept the licence terms. If you cannot accept these licence terms, please cancel the installation and contact Cambridge Research Systems Ltd.:
- On the next screen, select the location into which the software is to be installed, and whether to include desktop shortcuts or not.
- Click Next, then click Install. Wait for the installation to complete:
- After clicking Finish, the HardingFPA-X *Supervisor* will have been fully installed, and will appear (by default) under the HFP-A-X group on the Start Menu.

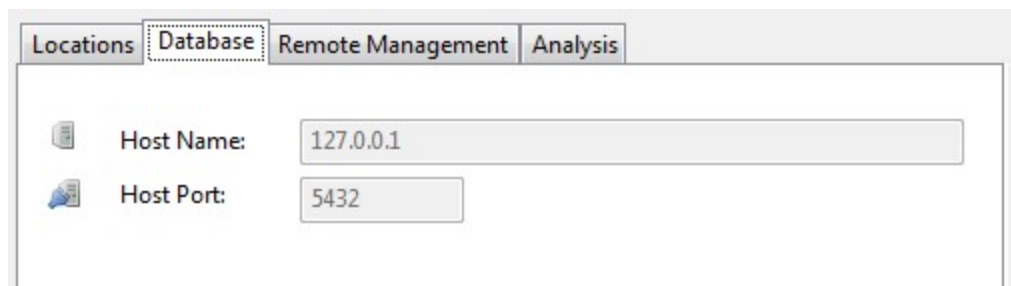
HardingFPA-X Components and Settings

When you have installed the applications to the required computers, each will need to be configured in order for the applications to know where to find the database, and which paths to use to find the movies, results files etc.

This section describes the operation and settings (configuration) for the HardingFPA-X System. Some settings are common between applications such as database settings, so please refer to the whole section if you cannot find the information you need.

Common Database Settings

Many of the applications in the HardingFPA-X System (*Submit*, *Monitor*, *Supervisor* and *Analyser*), have common settings which relate to the location on the network of the HardingFPA-X Database component. These settings can be found under *Tools* → *Settings* on these applications and appear as shown below.



Host Name: This is the host name or IP address of the computer on which the database is running. There must be only one HardingFPA-X Database on the network for a given system.

Host Port: This is the port number of the database in the system. Unless this was changed during installation, it should be left at its default value of 5432.

Any changes made to the database settings will require the relevant application to be restarted before the changes are made.

Language Support

All HardingFPA applications support multiple languages. In many of the applications, English and Japanese can be selected from the *Tools* → *Languages* menu. All applications will start up in the local language by default, providing the translation files are present.

Note: only English and Japanese languages are currently officially supported, although you may translate the applications yourself (see below).

To add a new language yourself, find the /languages/ subdirectory of the application bundle (on OS X) or program directory (on Windows) and make a copy of each of the *.po files, replacing 'en' with the two letter code for your locale (e.g. MainUnit.fr.po for French). You may then enter the strings after msgstr in this file, between the "" quotation marks. Ensure the *.po files are saved in the UTF-8 format (without BOM). An example is shown below for some menus translated into Japanese. these would be stored in xxxx.jp.po.

```
#: main:rssettings
msgid "Settings"
msgstr "設定"
#: main:rshelp
msgid "Help"
msgstr "ヘルプ"
```

Note that on the OS X platform, you must set one of the environment variables 'LC_ALL', 'LC_MESSAGES' or 'LANG' to the two letter code for your locale (e.g. 'jp' for Japanese etc.) for the applications to launch into your language automatically.

HardingFPA-X PostgreSQL Database

This is the main database which manages the parameters for all of the elements of the system, including jobs, file locations, licensed clients, HardingFPA-X system information etc. All of this information is managed transparently by the different components of the system. The database does not include analysis results themselves, only the locations of the files that contain them.

You must install and configure the database on one computer in your network so there will be only one Database installation to which all clients, applications and *Analysers* connect to. This database stores all information relevant to the jobs being analysed, pending jobs and completed jobs and allows the

HardingFPA-X system to be as flexible as possible (being extensible by allowing multiple clients and *Analyser* nodes).

The database does not require manual start-up. It runs as a service, and will therefore be running whenever the computer that it is installed on is switched on.

All management of the database is performed by the HardingFPA-X applications, and so there is no direct user interface to the database. It should be noted however that as the IP address for the database must be entered into the settings for most of the other HardingFPA-X applications, it is a requirement that the computer running the Database has a static IP address.

If required (for upgrade or maintenance for example), the database may be removed from the computer on which it was installed by:

Windows XP: using Add/Remove Programs to remove PostgreSQL.

Windows Vista/Windows 7: Using Programs and Features to remove PostgreSQL.

OS X: Run the script `uninstall.sh` which can be found in `/Library/HfpaPostgreSQL/bin/`

HardingFPA-X Launcher

This application is a small helper application that makes sure that the *Analyser* is always active and ready to receive jobs. It starts up automatically when the *Analyser* is started.

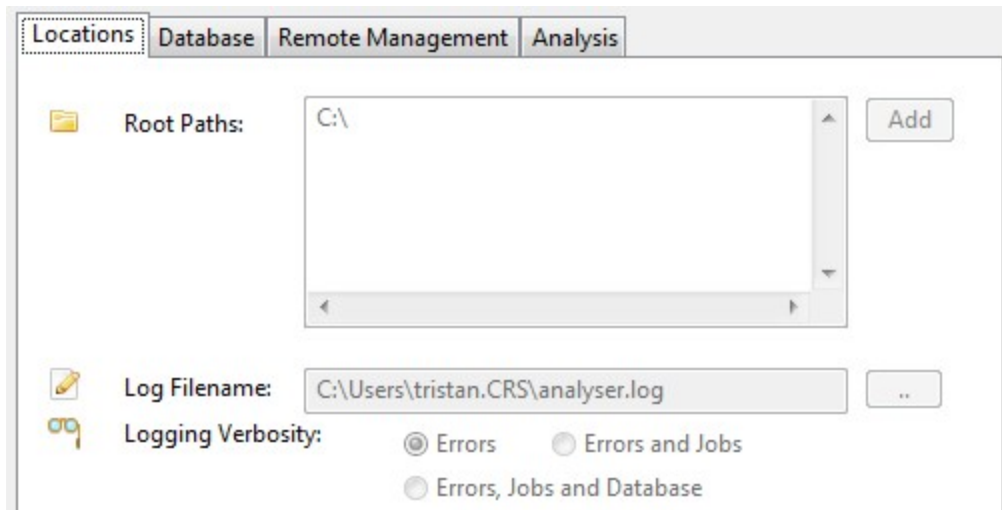


When the *Analyser* automatically shuts itself down to clear resources allocated by the movie frameworks (in the intervals specified in Maximum Job Length Analysis Rules), or if the *Analyser* is inadvertently shut down for any reason, the *Launcher* will restart it.

There is no configuration necessary for the *Launcher*. It will run in the system tray on Windows.

Storage and Paths

In order for the HardingFPA-X System to operate correctly, all computers must be able to access movie files that they would like to submit or analyse, and the applications must be correctly configured to allow this. The HardingFPA-X applications all provide a *Location Settings* section in order to specify this information.



In a typical usage scenario, the movie files are stored on a SAN or shared network drive. In this case, all of the *Analysers* and *Submits* in the system must be able to access this SAN, in order to be able to locate the files in question. The actual path used to access the same file from more than one computer may not be identical, so the Settings screen (*Tools* → *Settings*) for the *Analyser*, *Monitor* and *Submit* contain the following settings in order to specify the location of:

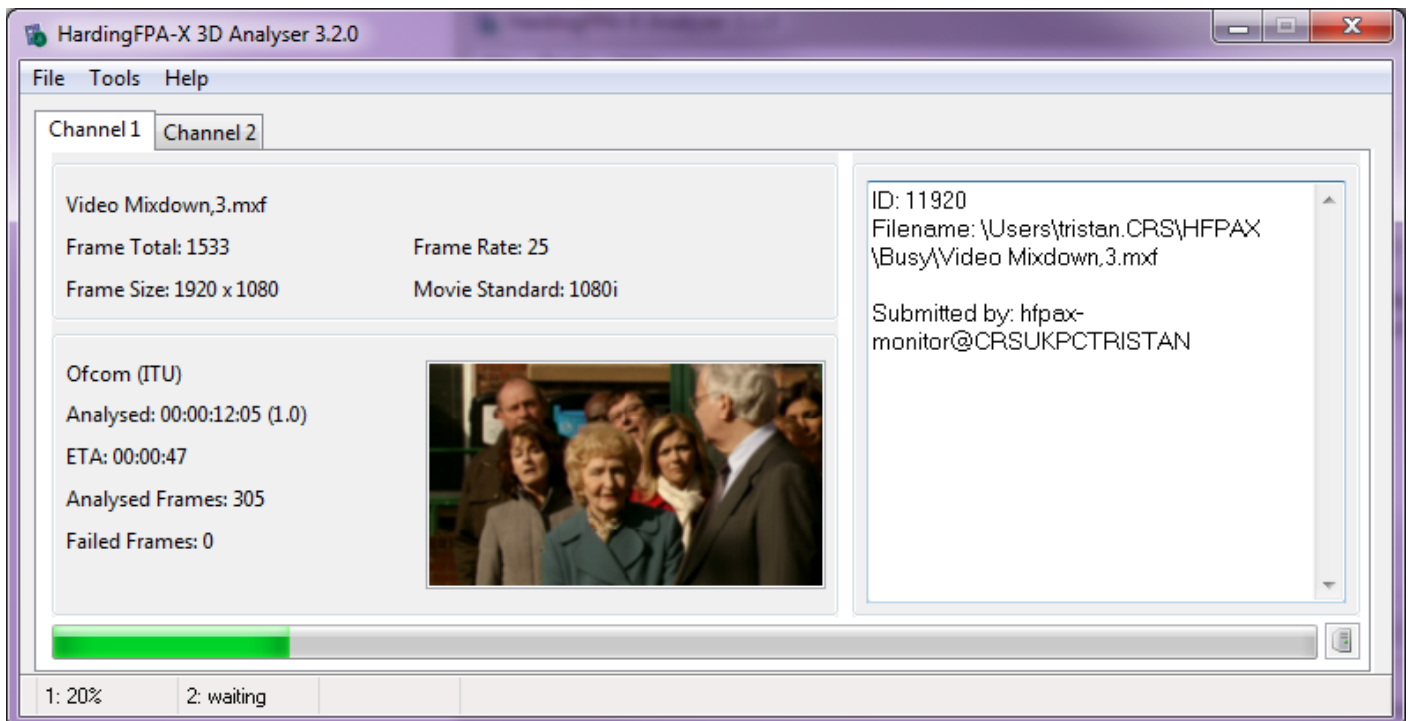
Root Path: This is the location of the root of your storage system. This folder will be the same actual folder for both the server computer and each of the clients, but as the folder might be mounted differently on each of the clients, this is the mount point that refers to that location from the **current computer** only, and so might be different on each computer. The *Analyser* allows for multiple root paths (on separate lines) to allow for more complex setups.

Log Filename: Some of the HardingFPA-X applications have an entry for *Log Filename*. In this box, enter the filename for the text file where you would like the application to store its logging information. This file may be of use if any problems occur, as the log can be checked for unusual circumstances. The *Log Filename* can be at any file location on the system which the application has write permissions to. The *Analyser* has the additional option of selecting a *Logging Verbosity* level (which can be seen on the

screenshot above). This allows the amount of information logged in the log file to be controlled. It may be useful to increase the *Logging Verbosity* if problems are experienced when running the *Analyser*.

HardingFPA-X Analyser

This is the main core of the HardingFPA-X functionality, and is a powerful multi-threaded Flash and Pattern Analyser based on the HardingFPA core with the ability to analyse multiple video clips simultaneously. Depending on your licence, your Analyser may support up to four channels simultaneously. If more than four channels are required, an additional *Analyser* node needs to be licensed and installed on another computer.



The Analyser graphical interface above shows the *Analyser* operating in two-channel mode, and therefore shows two identical tabs, one for each of the individual analyser channels, which detail the progress of the jobs currently being analysed. The information displayed on each of these tabs is explained below:

- The top left panel shows the name of the movie file being analysed, the total number of frames in the movie, the frame rate, frame size and movie standard.
- The bottom left panel shows the current state of the job for the selected analyser channel. The length of movie that has currently been analysed is shown, followed by a number in brackets, in the example above 1.0, which gives the current multiple of realtime being achieved by this *Analyser*

channel. For example 2.0 would mean that the movies on this *Analyser* core are being analysed two times faster than realtime. In the example, 1.0 refers to 1.0x realtime. ETA refers to the estimated time remaining before the job is complete and is only visible when it can be calculated, which may depend on the system and network load. Finally a count of the number of frames that have failed the guidelines is displayed under *Failed Frames*.

- The panel on the right shows the job ID number and filename of the job. Additionally, any errors or other information used for diagnosing problems is shown in this box. When the *Analyser* is first started, this panel will initially display the version information for the *Analyser*, and its analysis core (s).
- The status bar at the bottom of the window shows the progress of the jobs on each channel.

It is possible to have multiple multi-channel *Analysers* on the same network to spread the workload and provide greater throughput, provided that they are all configured to connect to the same central database.

SD Legacy Mode

The HardingFPA-X has a version 2.5 legacy analysis option for when comparisons with earlier HardingFPA analyses are necessary (i.e. those that have come from versions 1.x of the HardingFPA-X, and to have some level of compatibility with HardingFPA V2.5x SD-SDI Standalone tape-based systems).

If this mode is required it will need to be selected on **ALL** *Analysers* in your system.

To select SD Legacy Mode, select the Settings option on the Tools menu of the *Analyser*, check the box in the Analysis section to enable SD Legacy Mode, and then close and restart the *Analyser* application.

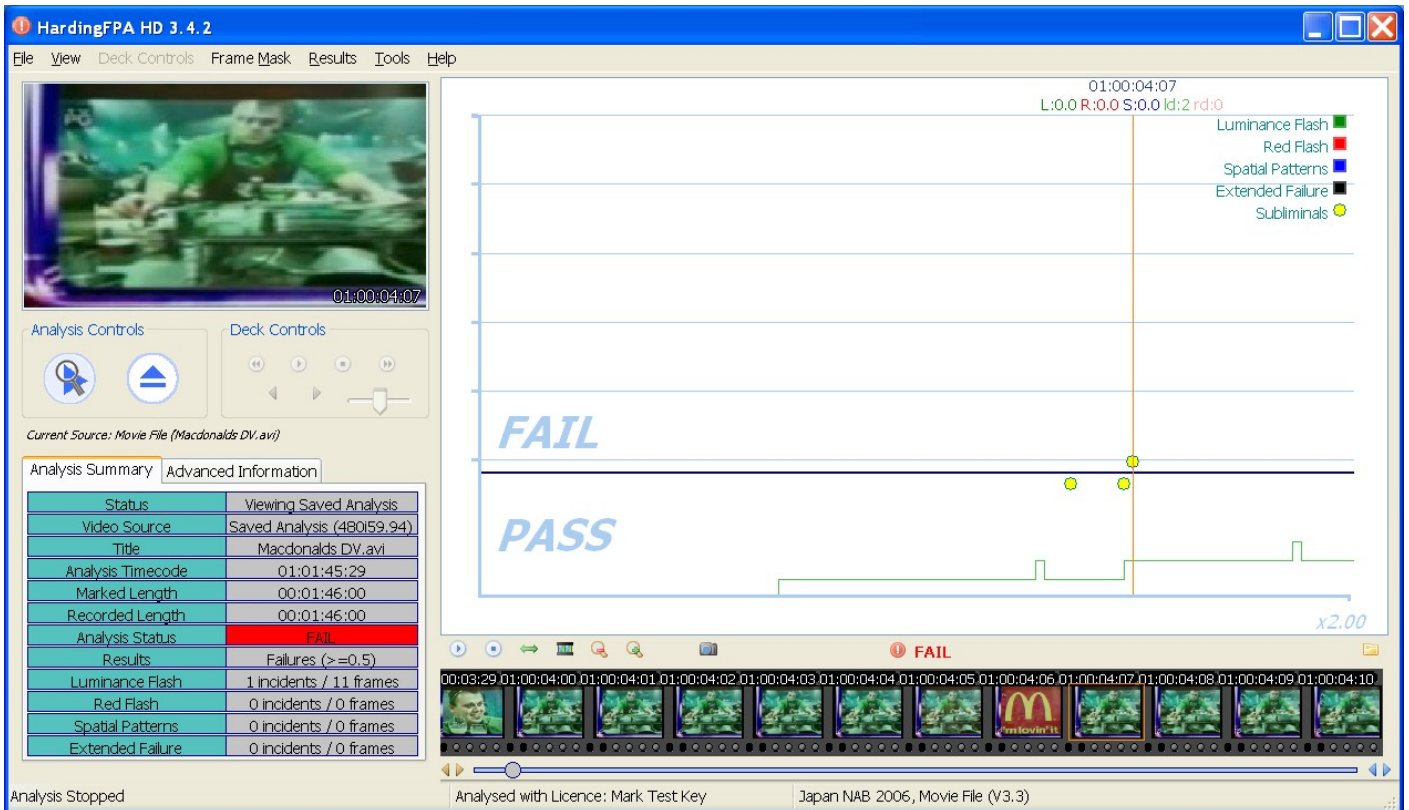
Please note, Legacy Mode is only available for Standard Definition (SD) video analysis. All HD formats will be analysed with Version 3 analysis algorithms. Once the application is restarted, the *Analyser* will only analyse in legacy mode unless the above change is reversed. The *Analyser* will indicate legacy mode analysis in the graph viewed by the HardingFPA-X Client applications and on any PDF results certificate, and also on the Job Information on the *Analyser*.

Subliminal Event Analysis

If enabled in the *Settings*, the system will test the incoming material for potential subliminal frames at the same time as it is analysing for the usual flash and pattern criteria.

These frames will be identified on the graph visualisation as yellow circles, shown in the example below. A yellow circle below the PASS/FAIL line indicates a subliminal warning, and one above the line indicates a subliminal failure.

The appearance of the subliminal results can be switched on and off in the usual manner, by clicking on the Subliminals yellow circle in the top-right corner of the graph.



What is a Subliminal Event?

The insertion of an image (or images) with contents different from the preceding or following images which is too short for the viewer to be consciously aware of. The duration of a subliminal event is user-configurable. The default settings are currently up to 0.05 seconds for a subliminal failure and up to 0.2 seconds for a subliminal warning. A subliminal event may be inserted at a scene change as it is not necessary for the images preceding and following the subliminal event to be similar. Additionally, a subliminal event may involve inserted text or a change to only part of the image.

It is important to note that a subliminal event must contain information. An inserted blank image that is black, white or uniform grey does not mean anything, and is therefore not regarded as subliminal.

Examples of subliminal and non-subliminal events

Example 1



This is a subliminal event because of the inserted images. The blue border in the 7th image shows the end of the subliminal event.

Example 2



This is a subliminal event because of the inserted picture of a person's head. The blue border in the 7th image shows the end of the subliminal event.

Example 3



This is a subliminal event because of the inserted text. The blue border in the 7th image shows the end of the subliminal event.

Example 4




This is not a subliminal event because the inserted images are just the inverse of the preceding and following images and therefore contain no new information.

Example 5



This is not a subliminal event because the inserted images contain no information.

Example 6



This is not a subliminal event because none of the images contain any information.

Example 7



This is a subliminal event because the inserted images contain information. The blue border in the 7th image shows the end of the subliminal event.

Example 8



This contains more than one subliminal event as shown by the blue border images.

How does the HardingFPA Detect Subliminal Events?

The HardingFPA searches for two sudden uncorrelated changes which occurred at different times but at the same position in the image sequence. Changes which are blended in over several frames are not considered as potentially subliminal as they would be consciously noticed by the viewer.

The HardingFPA can log the occurrences of two sudden uncorrelated changes and issue a subliminal event or a subliminal warning depending on its duration. These are indicated as yellow dots above and below the failure line respectively and are indications of where subliminal insertions may have occurred.

It is important to note that the HardingFPA may issue a large number of subliminal failures or warnings especially when analysing rapidly flashing imagery. This does not necessarily mean that the video contains a large number of subliminal insertions but merely that the HardingFPA has detected a large number of

potential candidates. It is ultimately up to the user to verify whether these flagged events constitute subliminal insertions or benign uncorrelated changes.

Why does the HardingFPA indicate a subliminal event one frame after it has occurred?

The HardingFPA needs to detect the start and end times of any potential subliminal insertion to determine whether the effect was rapid enough to be viewed subconsciously. Therefore, the HardingFPA can only log a subliminal event after it has disappeared.

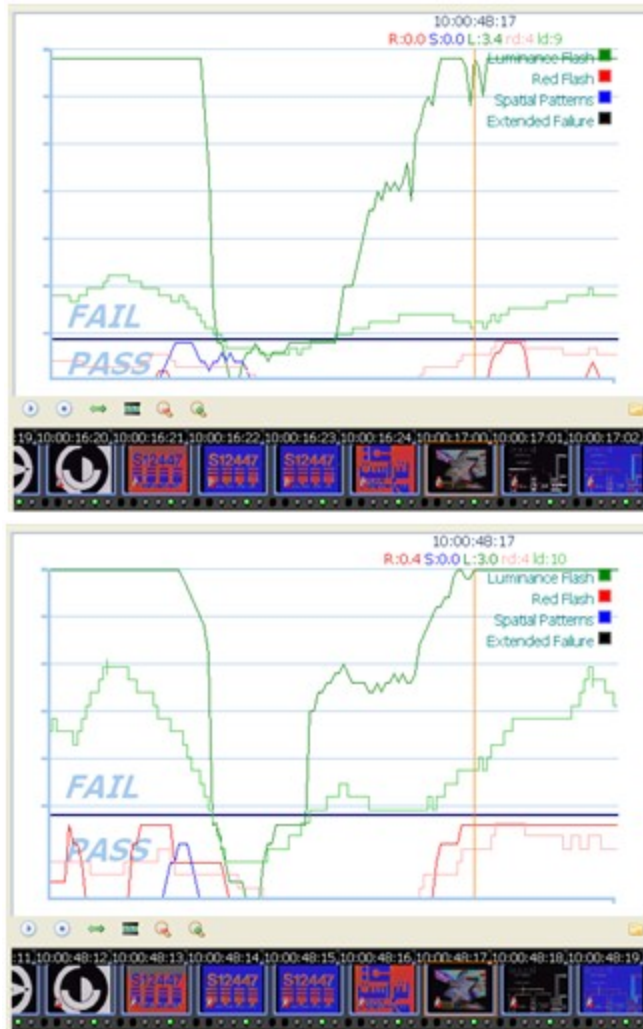
Version 3 Analysis Algorithms

HardingFPA software now features new analysis algorithms, which are better tuned to High Definition and File-based work.

The Version 3 analysis algorithms are better suited to accommodate subtle changes in the image data, and provide much closer results when testing the same material repurposed either into a different video format, or encoded with a different codec. The main differences between the legacy algorithms and the Version 3 algorithms are detailed below:

Different Graph Scaling

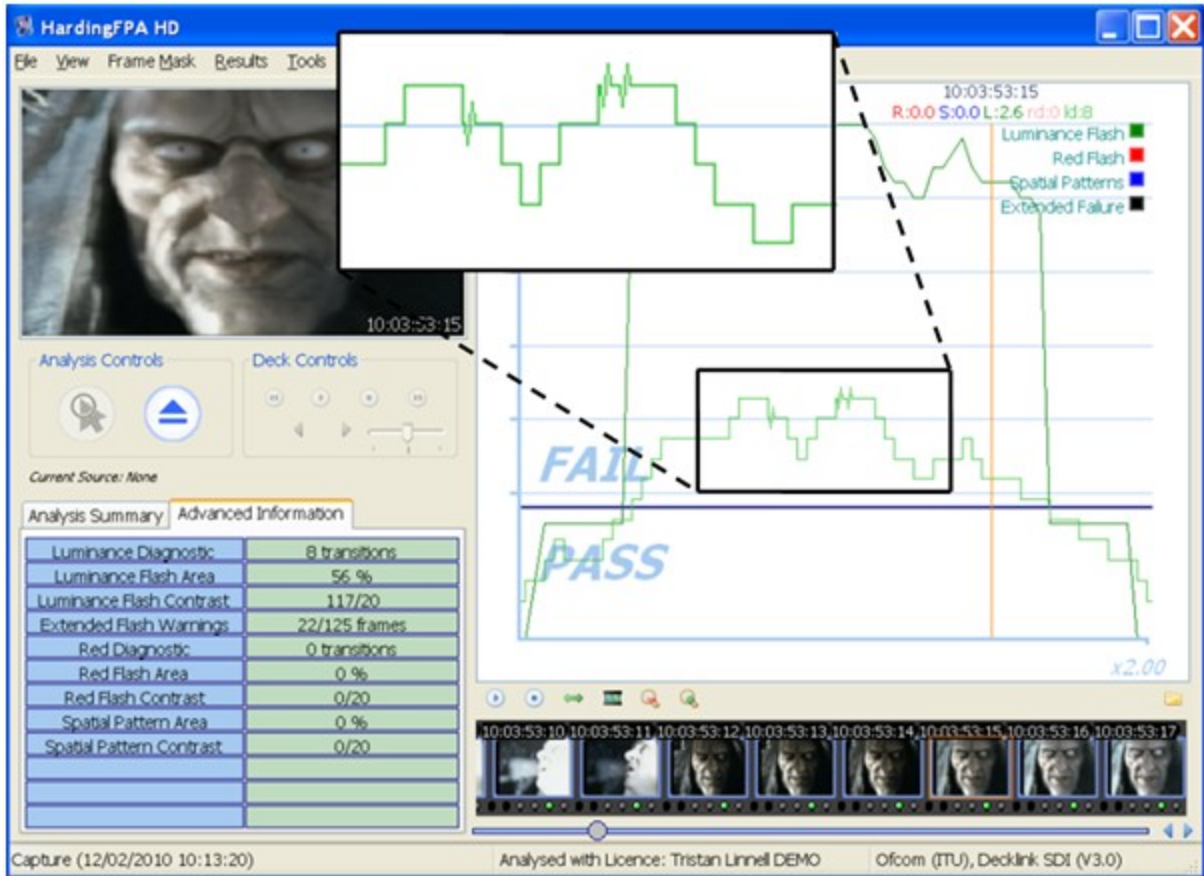
The HardingFPA generates risk values using the same range as its predecessor (i.e. 0 to 3.4) but displays the graphical data using a revised vertical scale. This modified scale allocates much more vertical space for risk trace warnings and diagnostic trace steps but only displays risk traces up to the value of 3.0. Risk traces values from 3.1 to 3.4 are still logged as part of the results files but are graphically displayed capped at 3.0.



(left) Version 2 and (right) Version 3 Graphing


The Squiggle

The HardingFPA gives enhanced visual diagnostics when an incoming transition coincides with an outgoing transition from one exactly second earlier. The Version 3 algorithms insert a squiggle (see below) to indicate when the diagnostic trace has simultaneously gained and lost a transition over the most recent second between video frames.




Analysis Results

The Version 3 algorithms give results which are broadly similar to those generated by version 2.5. The figure below shows the results of both versions when analysing the same video input under the same guidelines:



However, the results between the two versions will not be identical. The Version 3 algorithms will, in general, be more slightly more lenient to complex, rapid motion:




... but more strict to examples of powerful, localised flashing:

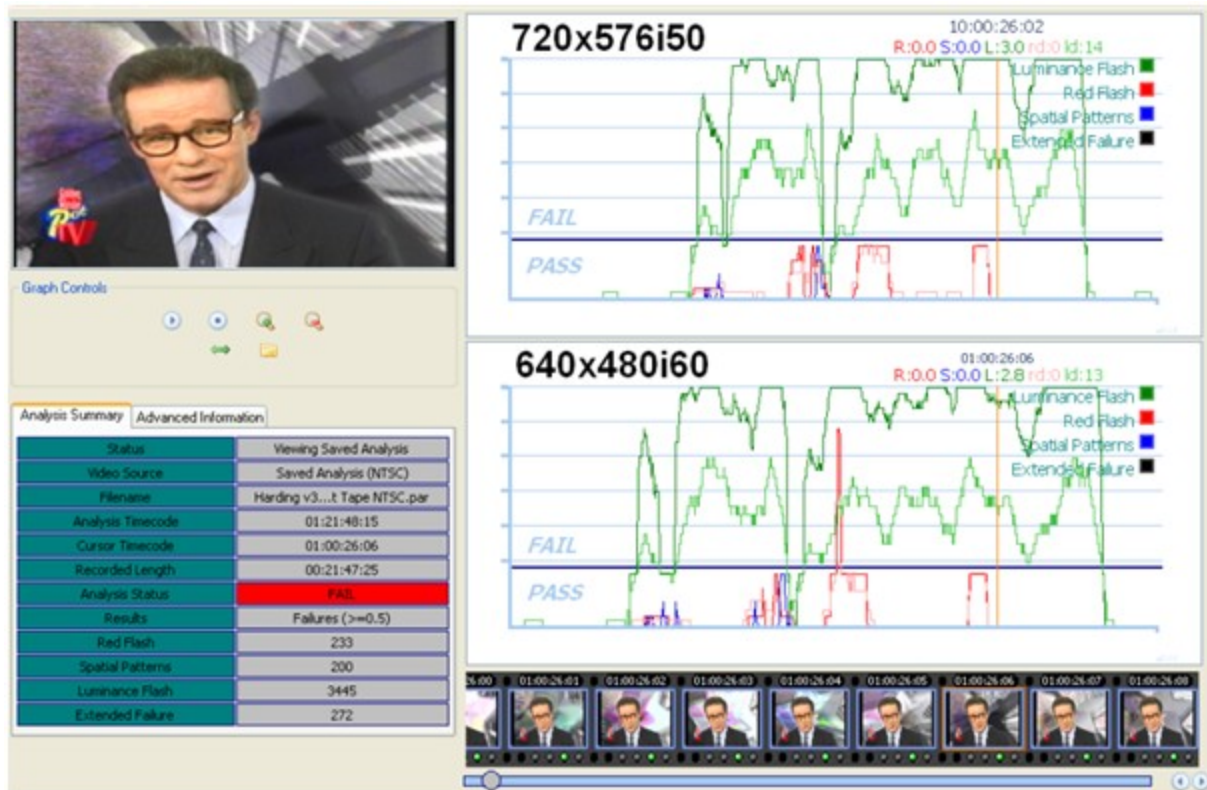


Most importantly of all, the Version 3 algorithms have been designed to be as format-agnostic as possible. Changes in file formats or codecs will alter the underlying video data even if these changes are not visually apparent. Here, the same video has been encoded at the same resolution using two different codecs. The absolute differences between the two images are shown in the third image as deviations from mid grey.






The Version 3 algorithms generate highly consistent results from different image resolutions and frame rates. Here, the same movie has been analysed in 720x576i50 and 640x480i60 formats. The only noticeable difference is in the horizontal graph scale due to the different frame rates.



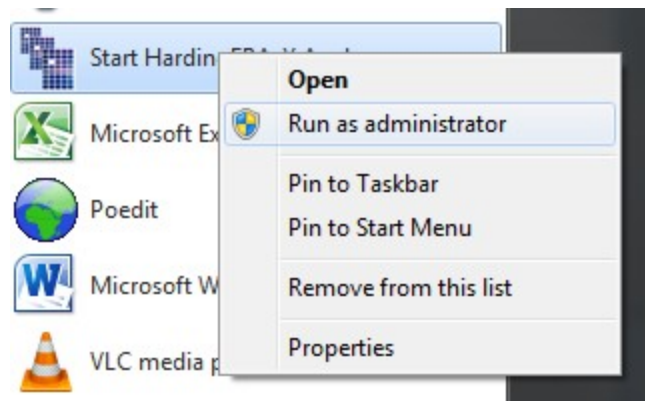
The graph below shows three sets of luminance flash risk results of the same movie analysed in SD-576i50, HD-720p50 and HD-1080i50 formats, and highlights the considerable similarities in the results:



Settings

This section details all of the settings in the HardingFPA-X Analyser.

If the settings are read only, then close and re-run the *Analyser* as an administrator user. In Windows, right-click the application icon and select *Run as administrator*, as shown below:

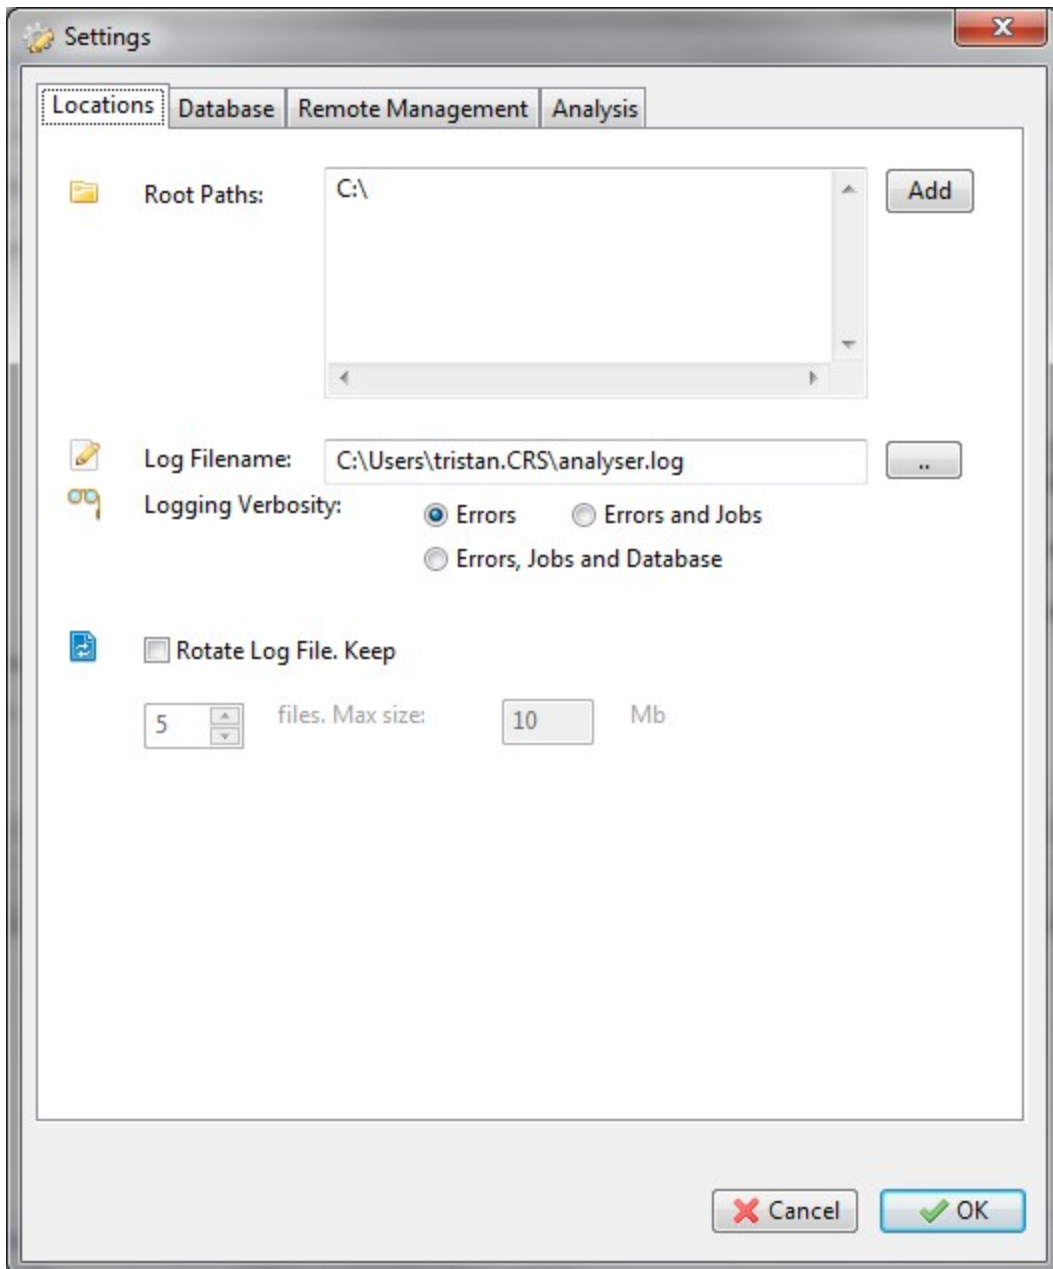


Locations

This screen details all of the settings relating to file paths and locations.

The *Analyser* can have multiple Root Paths (please see the section on paths for more information).

Material is then looked for in each of these root paths in the order in which they appear in this list.



The log verbosity can be set in the Settings screen to one of the following settings:

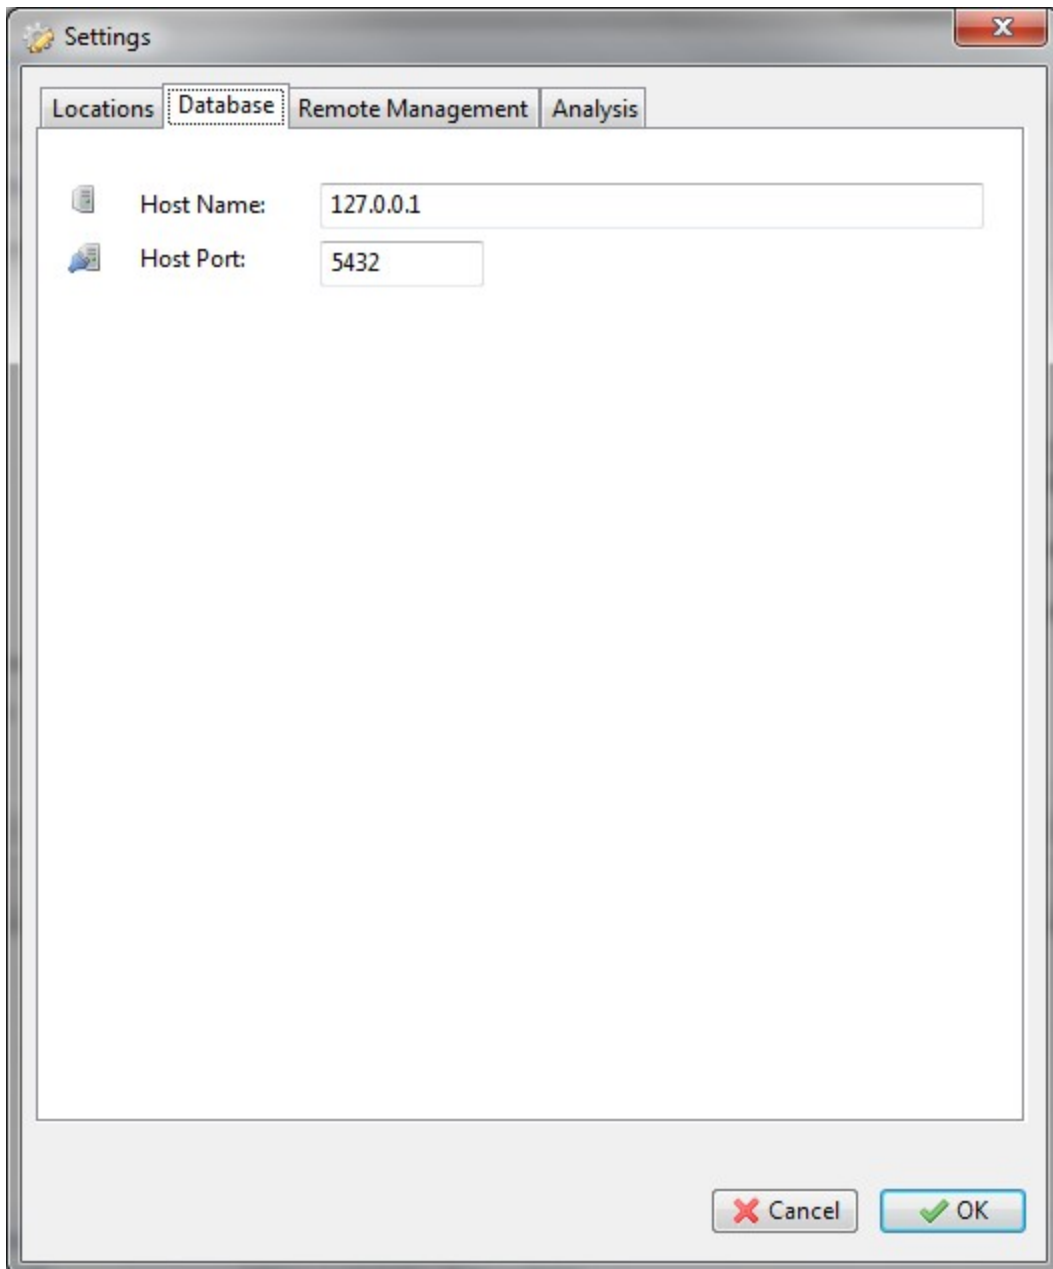
- *Errors*: Log errors only.
- *Errors and Jobs*: Log errors but add logging based on what each of the analyser channels is doing with the jobs (opening files, closing jobs, aborting jobs etc).
- *Errors, Jobs and Database*: The same as Errors and Jobs but with the addition of logging based on what communication is occurring with the database. This can be useful if unusual effects are

noticed that may be the result of problems with the database installation. Note that this setting creates a significantly larger log file than the other settings, because the *Analyser* has continuous communications with the database, so be sure there is enough free disk space for the log file.

If you would like to rotate the log file created, you may check *Rotate Log File* and select how many files to keep, and how big the file must get to before it is rotated. Rotated log files have numbers appended to their filenames.

Database

This screen details all of the settings relating to the Job Queue database.



Host Name: This is the host name or IP address of the computer on which the database is running. There must be only one HardingFPA-X Database on the network for a given system.

Host Port: This is the port number of the database in the system. Unless this was changed during installation, it should be left at its default value of 5432.

Any changes made to the database settings will require the Analyser to be restarted before the changes are made.

Remote Management

This screen details all of the settings relating to remote management.

The screenshot shows the 'Settings' window with the 'Remote Management' tab selected. The window is divided into three main sections: SNMP, Email, and SMTP Server.

SNMP Section:

- Enable SNMP Agent
- SNMP Port: 161
- Send SNMP Traps to: [Empty text box]

Email Section:

- Send PASS PDF results to:
- Send FAIL PDF results to:
- Send PDF (radio button)
- Send a link to PDF only
- Only send for jobs longer than: 1 minutes
- Email address list: tristan.linnell@crsltd.com

SMTP Server Section:

- Host Name: crsukmx01
- SMTP Port: 25
- Username: tristan
- Password: *****
- From Email Address: tristan.linnell@crsltd.com
- From Name: HardingFPA-X Analyser on CRSUKPCTRIST.
- Prefer TLS when connecting to the SMTP server

At the bottom of the window are 'Cancel' and 'OK' buttons.

Email

The email settings allow the *Analyser* to send the generated PDF certificate report (if settings allow it to be generated) to be sent via SMTP to a list of email addresses.

The checkboxes allow emails to be sent in the event of pass and/or fail results.

If the *Send a link to PDF only* radio button is selected, then the PDF itself will not be included in the email. Instead a link will be entered. Please be aware that these links will only work if the root paths are identical between the analyser computer and the device that the link is being opened on.

A list of destination email addresses can then be entered into the box beneath.

If *Only send for jobs longer than* is checked, then PDF certificates will only be emailed for movies that are longer than the specified length (in minutes).

SMTP Server

These are the settings for the SMTP server through which the emails are to be sent. Hostname, port, from address, from name, username and password are supported.


If the TLS checkbox is checked, an attempt to secure the connection with TLS will be made.

The password is stored in an encrypted form in the registry/settings file.

SNMP

Note that on the OS X platform, you may not be able to use the default SNMP listening port (161) unless the Analyser is being run as the root user. If the status icon (see below) shows an "error whilst binding", this is probably the cause. It is possible to change the port number in the settings.

The *Analyser* has the option to enable a Simple Network Management Protocol agent (SNMP server) so that you may monitor its operation through a third party SNMP manager. The settings for this can be found on the Settings window of the Analyser. Here the SNMP Agent can be turned on or off with the checkbox, the port number to listen on (UDP) can be selected, and (optionally), an IP address or hostname of a network manager to send SNMP Traps to in the event of any errors. All traps are sent to the default port **162**.



When the SNMP Agent is active and functioning correctly, the server icon is shown in the bottom right corner of the Analyser window, as shown below. In the event of any errors or changes of state, this icon will change. Hover the mouse over the icon to see the status of the SNMP Agent in text.



MIB:

The Analyser supports the following OIDs from the standard management MIB (.1.3.6.1.2.1.1):

- .1.3.6.1.2.1.1.1.0 - sysDescr. The name and operating system of the Analyser.
- .1.3.6.1.2.1.1.2.0 - sysObjectID The OID for HFPA-X (This is always .1.3.6.1.4.1.33056.1).
- .1.3.6.1.2.1.1.3.0 - sysUpTimeInstance The number of 100ths of a second that the SNMP Agent has been active in the Analyser.
- .1.3.6.1.2.1.1.5.0 - sysName The host name of the computer the Analyser is running on.
- .1.3.6.1.2.1.1.7.0 - sysServices The sysServices value for the SNMP layer (64).

The HardingFPA-X Products also have a private MIB (the .mib file for which can be found with the installation media, and is reprinted below) to facilitate HardingFPA-X specific management variables and traps. The root of these can be found at:

.1.3.6.1.4.1.33056.1

Please see the MIB for details on what each OID can be used for.

Please note that traps are sent with the community string "HardingFPA-X" in order to facilitate filtering of traps in your management client software.

The HardingFPA-X .mib is duplicated below, in ASN.1:

```
CRS-MIB DEFINITIONS ::= BEGIN

IMPORTS

OBJECT-TYPE

FROM RFC-1212
```



```
TRAP-TYPE

FROM RFC-1215;

cambridgeResearchSystems OBJECT IDENTIFIER

-- 1.3.6.1.4.1.33056

::= { 1 3 6 1 4 1 33056 }

hfpaX OBJECT IDENTIFIER

-- 1.3.6.1.4.1.33056.1

::= { cambridgeResearchSystems 1 }

hfpaXNotifications OBJECT IDENTIFIER

-- 1.3.6.1.4.1.33056.1.1

::= { hfpaX 1 }

hfpaAnalyserJobErrorTraps OBJECT IDENTIFIER

-- 1.3.6.1.4.1.33056.1.1.1

::= { hfpaXNotifications 1 }

hfpaAnalyserJobErrorGeneral TRAP-TYPE

ENTERPRISE hfpaAnalyserJobErrorTraps

VARIABLES {

hfpaAnalyserErrorChannel,

hfpaAnalyserErrorString}

DESCRIPTION

"A DirectShow error occurred during job processing. Error

string and channel number included"

::= 5500

hfpaXObjects OBJECT IDENTIFIER

-- 1.3.6.1.4.1.33056.1.2
```

```
::= { hfpaX 2 }

hfpaXStats OBJECT IDENTIFIER

-- 1.3.6.1.4.1.33056.1.2.1

::= { hfpaXObjects 1 }

hfpaAnalyserStats OBJECT IDENTIFIER

-- 1.3.6.1.4.1.33056.1.2.1.1

::= { hfpaXStats 1 }

hfpaAnalyserchannel1Filename OBJECT-TYPE

SYNTAX OCTET STRING

ACCESS read-only

STATUS optional

DESCRIPTION

"The filename of the movie file currently being analysed by
Channel 1"

-- 1.3.6.1.4.1.33056.1.2.1.1.1

::= { hfpaAnalyserStats 1 }

hfpaAnalyserchannel2Filename OBJECT-TYPE

SYNTAX OCTET STRING

ACCESS read-only

STATUS optional

DESCRIPTION

"The filename of the movie file currently being analysed by
Channel 2"

-- 1.3.6.1.4.1.33056.1.2.1.1.2

::= { hfpaAnalyserStats 2 }
```

hfpaAnalyserchannel3Filename OBJECT-TYPE

SYNTAX OCTET STRING

ACCESS read-only

STATUS optional

DESCRIPTION

"The filename of the movie file currently being analysed by
Channel 3"

-- 1.3.6.1.4.1.33056.1.2.1.1.3

::= { hfpaAnalyserStats 3 }

hfpaAnalyserchannel4Filename OBJECT-TYPE

SYNTAX OCTET STRING

ACCESS read-only

STATUS optional

DESCRIPTION

"The filename of the movie file currently being analysed by
Channel 4"

-- 1.3.6.1.4.1.33056.1.2.1.1.4

::= { hfpaAnalyserStats 4 }

hfpaAnalyserchannellPercent OBJECT-TYPE

SYNTAX INTEGER (0..100)

ACCESS read-only

STATUS optional

DESCRIPTION

"The percentage complete value for the job currently
on channel 1"

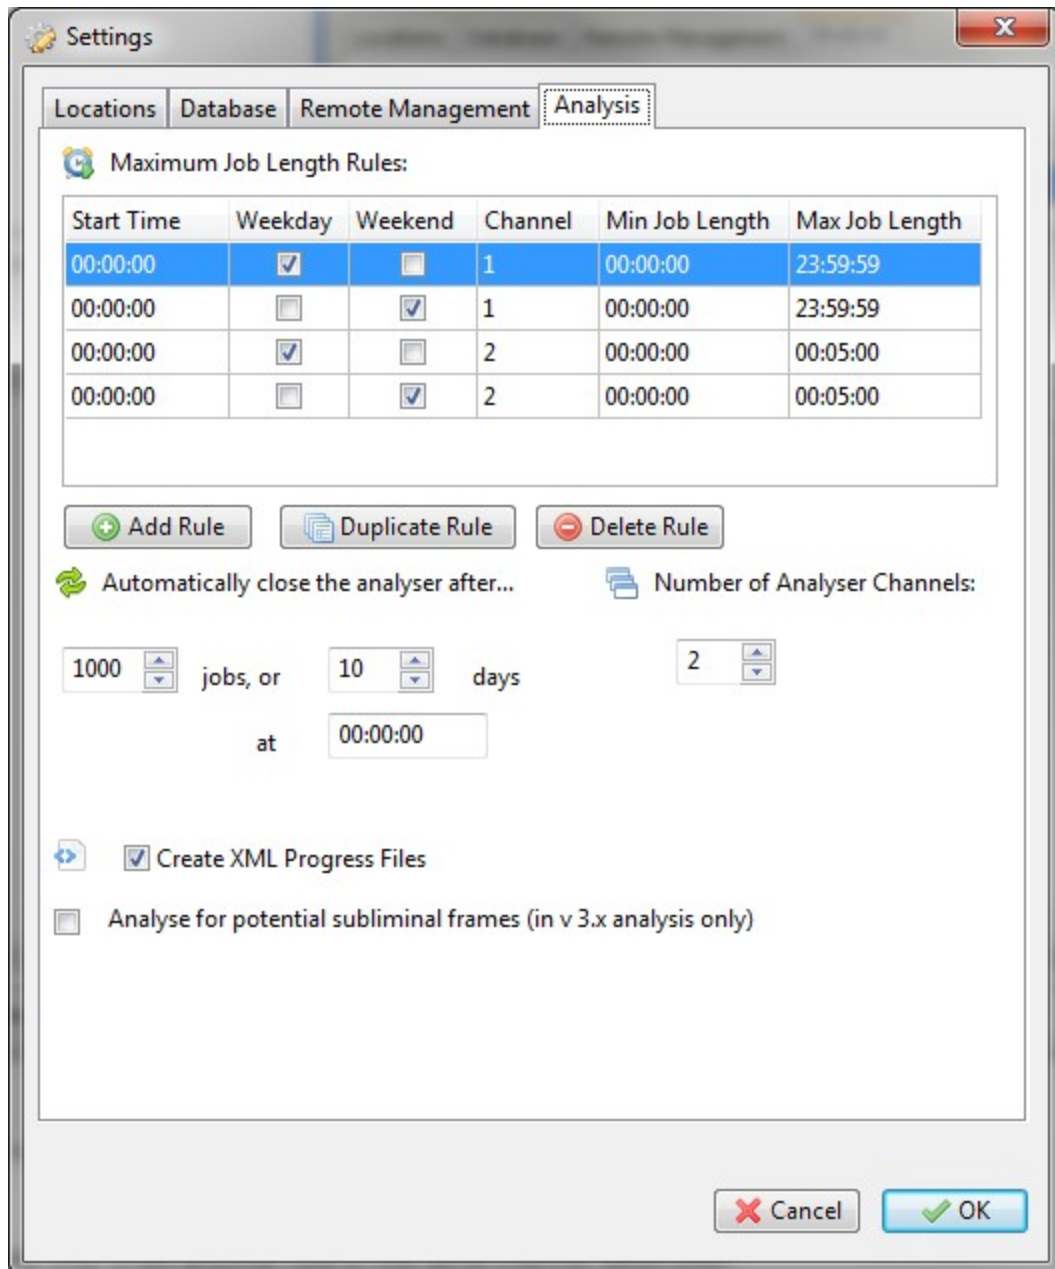
```
-- 1.3.6.1.4.1.33056.1.2.1.1.5
 ::= { hfpaAnalyserStats 5 }
hfpaAnalyserchannel2Percent OBJECT-TYPE
SYNTAX INTEGER (0..100)
ACCESS read-only
STATUS optional
DESCRIPTION
"The percentage complete value for the job currently
on channel 2"
-- 1.3.6.1.4.1.33056.1.2.1.1.6
 ::= { hfpaAnalyserStats 6 }
hfpaAnalyserchannel3Percent OBJECT-TYPE
SYNTAX INTEGER (0..100)
ACCESS read-only
STATUS optional
DESCRIPTION
"The percentage complete value for the job currently
on channel 3"
-- 1.3.6.1.4.1.33056.1.2.1.1.7
 ::= { hfpaAnalyserStats 7 }
hfpaAnalyserchannel4Percent OBJECT-TYPE
SYNTAX INTEGER (0..100)
ACCESS read-only
STATUS optional
DESCRIPTION
```

```
"The percentage complete value for the job currently
on channel 4"
-- 1.3.6.1.4.1.33056.1.2.1.1.8
 ::= { hfpaAnalyserStats 8 }
hfpaAnalyserErrors OBJECT IDENTIFIER
-- 1.3.6.1.4.1.33056.1.2.1.1.5000
 ::= { hfpaAnalyserStats 5000 }
hfpaAnalyserErrorString OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS optional
DESCRIPTION
"The last error string that occurred. Also sent
with error traps"
-- 1.3.6.1.4.1.33056.1.2.1.1.5000.4
 ::= { hfpaAnalyserErrors 2 }
hfpaAnalyserErrorChannel OBJECT-TYPE
SYNTAX INTEGER (1..4)
ACCESS read-only
STATUS optional
DESCRIPTION
"The channel number that the last error occurred on. Also sent
with error traps"
-- 1.3.6.1.4.1.33056.1.2.1.1.5000.1
 ::= { hfpaAnalyserErrors 1 }
```

END

Analysis

This screen details all of the settings relating to analysis of movies.



Maximum Job Length Analysis Rules:

These settings allow fine control over when each channel is allowed to run jobs of a certain length. As many rules as are required can be added to the list.

- **Start Time:** This is when the rule starts. It will only stop when another rule of the same channel takes over.
- **Weekday:** If the rule should apply on a weekday (Monday - Friday).
- **Weekend:** If the rule should apply on a weekend (Saturday - Sunday).
- **Channel:** Which channel to apply the rule to.
- **Min Job Length:** The minimum length of job that the selected channel will accept.
- **Max Job Length:** The maximum length of job that the selected channel will accept.

Automatically close the analyser after:

This controls when the *Analyser* shuts itself down, in order to be restarted by the *Launcher*. The settings are in number of jobs, and number of days (Also specify the time of day to automatically close). The *Analyser* will close after whichever occurs first.

The checkbox Create XML Progress Files can be unchecked if you do not want the *Analyser* to create XML Progress Files during the analysis of jobs.

This is also where the [SD Legacy Mode](#) and [Subliminal Event Analysis](#) can be turned on or off.

Results Formats

The *Analyser* is able to produce results of different types, depending on your usage scenario. With the use of the *Monitor* application or XML token file job submission, the different types of results files can also be placed in different folders. The available results types are explained in this section.

PAP/PAR/PAM Results

These are proprietary CRS HardingFPA Results files, which contain numerical (PAR), pictorial (PAP) and Frame Mask (PAM) details of the analysis outcome. These can be viewed using the HardingFPA-X Viewer application.

PDF Results

A PDF certificate showing the overall pass or fail status. Fail reports can also be set up to provide additional diagnostic information, including the identification of failure regions with thumbnails. The PDF format is the same as those obtained from other HardingFPA applications.

CSV Results

Comma Separated Variables of the raw data from the analysis. This can be used for more advanced set-ups, should you require access to the raw data relating to the results traces. The file is presented as follows:

```
"HH:MM:SS:FF", "Flash", "Red", "Spatial", "Extended", "FlashDiag",  
"RedDiag"  
  
"00:00:00:00", 0, 0, 0, 0, 0, 0  
"00:00:00:01", 0, 0, 0, 0, 0, 0  
"00:00:00:02", 0, 0, 0, 0, 0, 1  
"00:00:00:03", 0, 0, 0, 0, 0, 1  
"00:00:00:04", 0, 0, 0, 0, 0, 1  
"00:00:00:05", 0, 0, 0, 0, 0, 0  
"00:00:00:06", 0, 0, 0, 0, 2, 1  
"00:00:00:07", 0, 0, 1, 0, 2, 0  
"00:00:00:08", 0, 0, 2, 0, 2, 2  
"00:00:00:09", 0, 0, 2, 0, 2, 1  
  
... etc.
```

XML Results

This is another way of obtaining the raw results data for the analysis, this time in eXtensible Markup Language. This file contains more information than the CSV format. All of the fields available from the

PAR file (viewed using the Viewer) are stored in the XML results file. Please refer to the HardingFPA-X System User's Manual for detailed information on the result fields. The layout of this file is as follows:

```
<?xml version="1.0" ?>

<fpasvc>

  <header>

    <fpasvc version="HardingFPA-X Analyser 3.1.0" />

    <file name="test.mov" assetnumber="as12345" />

    <video framerate="25.00" width="720" height="576" />

    <analysis totalframes="00:00:40:24" />

    <analysis endframe="00:00:40:24" />

  </header>

  <summary analyser="HardingFPAX" version="2115" analysis_
standard="OfCOM (ITC)" mode="PAL" scan="Interlaced"
format="UYUV" status="Failed" error="">

    <item type="luminanceflash" count="54"
warnings="559"/>

    <item type="redflash" count="5" warnings="57"/>

    <item type="spatialpattern" count="0" warnings="7"/
/>

    <item type="extended" count="0" />

  </summary>

  <results>

    <failsequence startframe="00:00:10:19">

      <failure timecode="00:00:10:19" flash="0.5"
red="0.0" spatial="0.0" extended="0">

    </failsequence>

  </results>

</fpasvc>
```

```
reddiag="7" flashdiag="7"/>  
  
<failure timecode="00:00:10:20" flash="0.7"  
red="0.0" spatial="0.0" extended="0"  
reddiag="7" flashdiag="7"/>  
  
<warning timecode="00:00:10:21" flash="0.9"  
red="0.0" spatial="0.0" extended="0"  
reddiag="7" flashdiag="7"/>  
  
<warning timecode="00:00:10:22" flash="1.1"  
red="0.0" spatial="0.0" extended="0"  
reddiag="7" flashdiag="7"/>
```

... etc, for each of the data points in the current failure/warning region sequence.

```
</failsequence>
```

```
<failsequence startframe="00:00:20:05">
```

... Repeats as before. One section of this type for each failure/warning region sequence...

```
</failsequence>
```

```
</results>
```

```
</fpasvc>
```

The *status* field in the summary section can take the following values:

- "Failed"
- "Passed"
- "Stopped by error XXYYZZAABB"

Where XX, YY, ZZ and AA are error codes:

XX : MXF Error code

YY : DirectShow Error Code

ZZ : QuickTime Error Code

AA: FFmpeg Error Code

BB: FPA Core Error Code

Note that warnings will only appear in the XML Results files if this option was selected when the job was submitted (using the *Monitor* application).

XML Progress Files

As jobs are being analysed, their progress can be tracked in an external automation system by reading from the progress *file*. In order to allow for the creation of these files, this option must be enabled in the Settings screen of the HardingFPA-X Analyser. The file name will consist of the Job ID, followed by the movie filename, and with the string *'_progress'* appended to the name. The contents of the file are in the following format:

```
<?xml version="1.0" ?>

<hfpax>

  <header>

    <analyser version="3.1.0" />

    <file name="/Users/tristan/hfpax/busy/testsequence2_
xvid.avi" assetnumber="as12345" />

    <video framerate="25.00" width="720" height="576" />

    <analysis startframe="00:00:00:00"
endframe="00:02:14:13" />

  </header>

  <status value="complete" position="00:02:14:13" progress="100"
/>

</hfpax>
```

The *status tag* in the example above indicates that the file has been analysed. The *progress* attribute indicates the percentage of the file that has been analysed.

The possible values of status are:

1. *busy*, the file is being analysed.
2. *complete*, analysis has successfully been completed. This does not indicate that the movie has passed.
3. *aborted*, the file analysis has been aborted.
4. *error*, there was an error during the analysis.

The position attribute in the status tag details the position of the analysis at the current time, or at the position of the error in the case of the error state. In the case of the error status, an additional tag is included with the error details e.g. `< error code ="4" message="QT ERROR 4: Visual context error" />`, which appears under the status tag. Error numbers below 100 indicate there was a problem with the movie file, and numbers above 100 indicate that there was an error with the analysis engine.

Other information contained within the progress file includes the analyser version, filename, frame rate, frame width and height etc. as well as an optional assetnumber field, which is specified if the job was submitted using an XML Token File with an Asset Number specified.

HardingFPA-X Monitor

This application is required if you require the use of Watch Folders or XML Token Files to submit jobs to the system. To use Watch Folders, a user simply moves or copies files to be analysed into a pre-designated input folder. The *Monitor* then identifies the job, moves it into a busy folder and submits the job to the *Analyser*. Completed jobs are automatically moved to the relevant folder after analysis. Different folders can be specified for jobs that pass, fail or cause an error.

This process can be easily automated in a more complex set-up. Options exist to move jobs to different folders depending on their pass/fail status, and support is given for multiple watch folders via wildcard search patterns or lists of folders. The XML Token File is an XML file which can be used to specify the parameters for a specific job, which are interpreted by the Monitor which then schedules the corresponding job.

The *Monitor* will submit movies as long as it has read/write permissions to them, they are greater than 2kb in size and they have one of the following file extensions:

mov, dv, avi, wmv, vob, mpg, mpeg, xml, mxf, mp4, m2v, flv, ts, qt, ps, 3gp, mkv, m2ts.

Note however that even if the *Monitor* successfully submits the job, the *Analyser(s)* in the system must have the appropriate codec correctly installed and configured in order to be able to analyse it, and must be able to locate the movie file on the network.


Tabs

The *Monitor* application features the following tabs:

The *Job Queue* and *Completed Jobs* tabs show the current jobs being analysed and waiting (*Job Queue*) and the jobs that have been analysed in the last 14 days (*Completed Jobs*). Please note that it is not possible to manipulate the *Job Queue* or view results via the *Monitor* application. The tabs here are for information only.

The *Folders* tab shows a list of all folders currently being watched by the *Monitor*. Movies that are copied or moved into these folders will automatically be entered into the job queue for analysis as soon as possible.

The *Logging* tab shows a log of activity performed by the *Monitor* including adding watch folders, queuing jobs, re-queuing jobs and any errors that occur. If any problems occur with the watch folder functionality, remember to check here for any unusual error messages.





Watch Folder Settings

This section defines the locations of the various folders used in the HardingFPA-X system, including the Watch Folders that the *Monitor* application collects jobs from.

The process is that movie files that are placed in the Movie Input folder are moved by the *Monitor* to the *Movie Busy* folder as soon as they are placed into the job queue, and are then moved to either *Movie Passed*, *Movie Failed* or *Movie Error*, depending on whether they pass or fail the analysis or cause an error during movie processing. The different results types are then placed in the relevant *Results* folders.

Multiple Watch Folder Sets

The Monitor application supports as many distinct watch folder setups as you like. Each one will act independently and have its own root path and set of settings. This can be useful to have different watch folder configurations for different analysis standards for example. Click the plus button () to add a new set of watch folders, and the minus () to remove the currently selected set.

To use multiple watch folders, please make sure that the Analyser has each of the Root Paths configured.

Standard Watch Folder Operation Mode

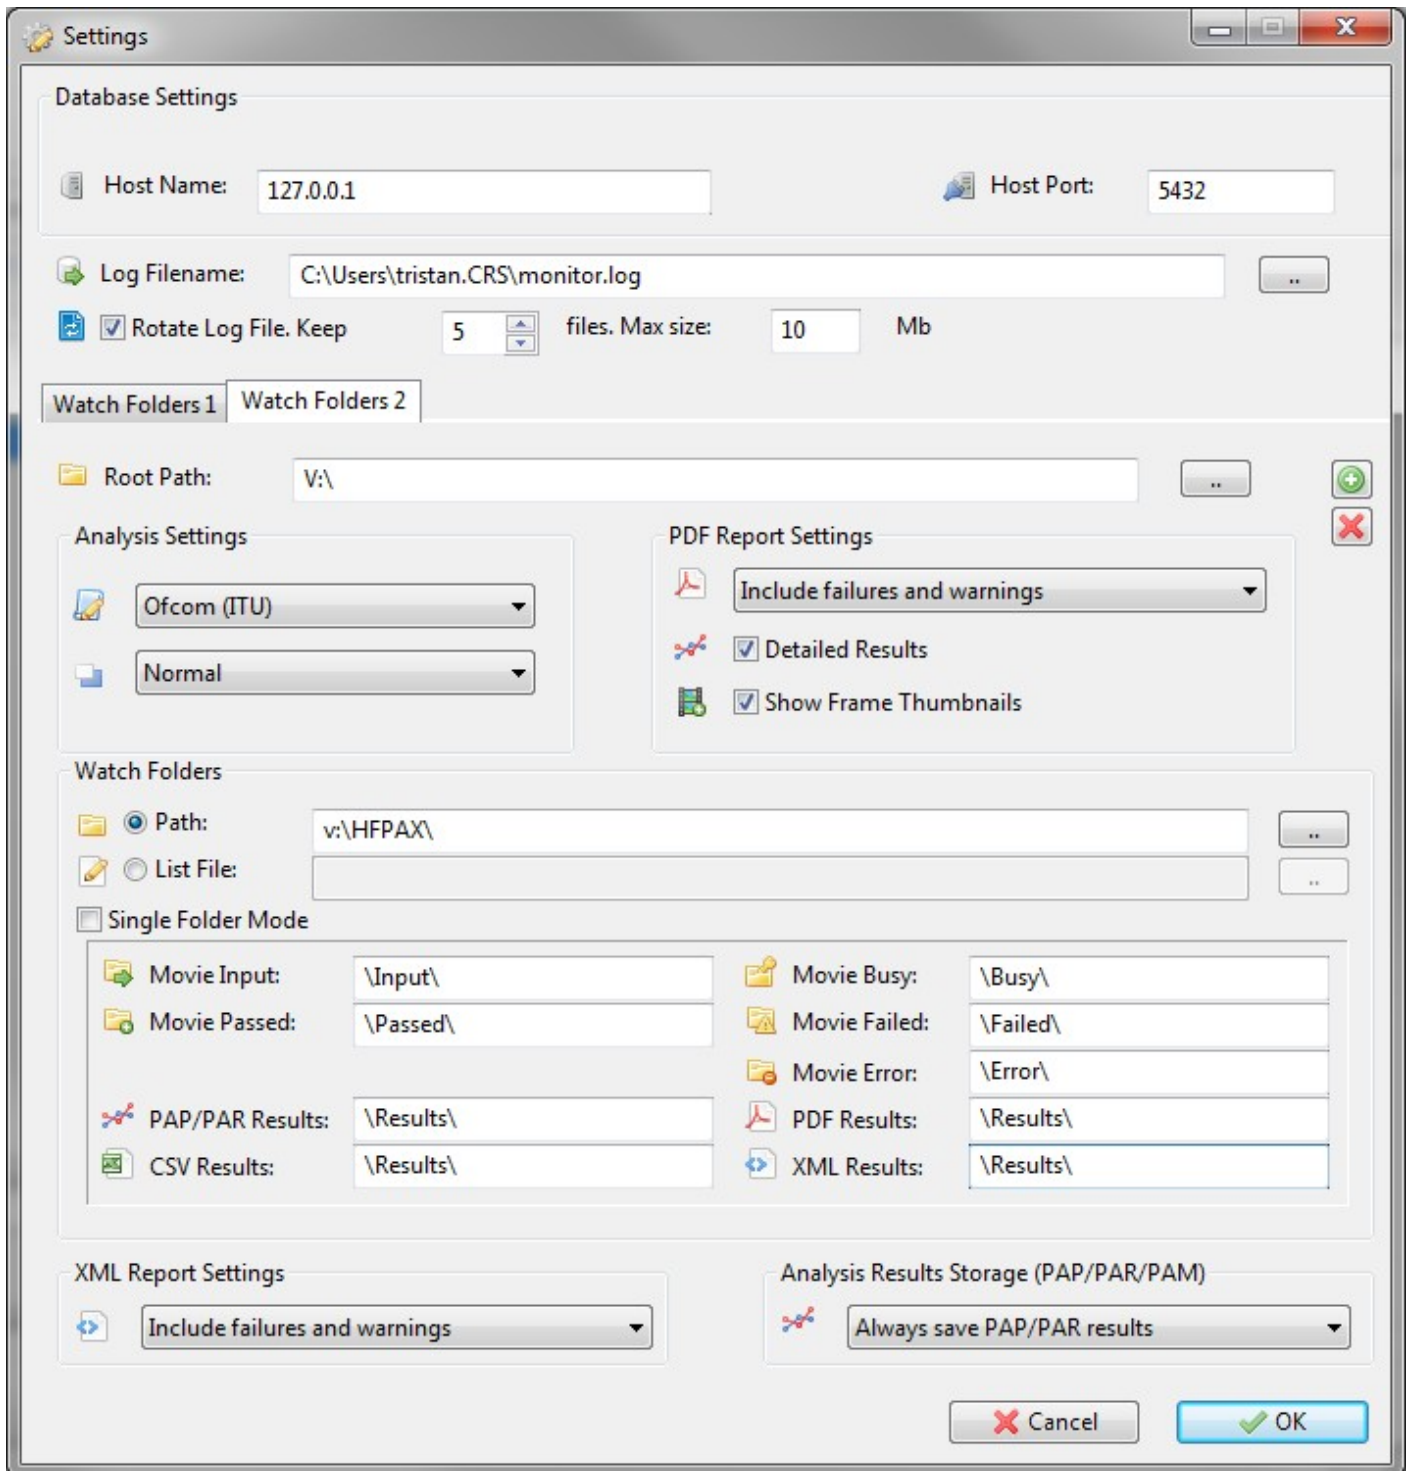
These instructions are for the conventional watch folder mode, where material is to be dropped/copied to an input folder, then moved to a busy folder during analysis, and later moved to output folders based on analysis result. These settings are duplicated for each watch folder set.

Path: This is the root folder, underneath which all of the following sub-folders are to be found. It is also possible to include wildcards in this path in order to have multiple watched folders.

Note that all of these multiple watch folders must have the same subfolder structure underneath them.

For example if the Path was set to /movies/test*, then the application will operate on the input, busy and output subfolders of all of the folders that match the wildcard, in this case the folders that start with /movies/test. For example /movies/test1, /movies/test2, /movies/test3 etc, all of which should be configured to contain the same pattern of subfolders for input, output and results.

List File: This is another way of specifying multiple watched folders. This path should specify a text file which contains each watch folder path on its own line of the file. The *Monitor* will then add these folders to the list of watch folders. Note the *Monitor* will expect to find the same subfolder structure (input, busy, output etc.) in each of the watch folders.



The purpose of each subfolder is as follows:

Movie Input: Where the *Monitor* should look for movie files to be analysed. Movie files can be moved or copied into this folder by users who wish to have the files analysed by the HardingFPA-X System.

Movie Passed: Where the *Monitor* moves any movie files that pass (after analysis is complete).

PAR Results: Where the PAP/PAR/PAM results files (that are used by the Viewer application) should be saved to when the analysis is complete.

CSV Results: Where the CSV results files are saved to after the analysis is complete.

Movie Busy: Where the *Monitor* moves the movie files to after they are added to the job queue.

Movie Failed: Where the *Monitor* moves any movie files that fail (after analysis is complete).

Movie Error: Where the *Monitor* moves any movie files that wither cause an error, resulting in abnormal termination of analysis, or are aborted by a supervisor.

PDF Results: Where the PDF results files are to be saved after the analysis is complete.

XML Results: Where the XML results files are to be saved after the analysis is complete.

In the *Analysis Standard* section, the analysis standard which all jobs submitted by this *Monitor* application will be analysed as can be set. Please note that the *Analysers* that pick up these jobs need to be licensed to analyse the standard that is chosen; therefore the HASP key in use on the *Analyser* computer must have the relevant guidelines licensed. Please contact Cambridge Research Systems if you require additional analysis standards.

Single Watch Folder Operation Mode

If required, it is possible to put one or more watch folder sets into *Single Folder Operation Mode*, whereby all assets including source files to be analysed and results files (PAP, PAR, CSV, XML etc.) are all stored in the same watch folder.

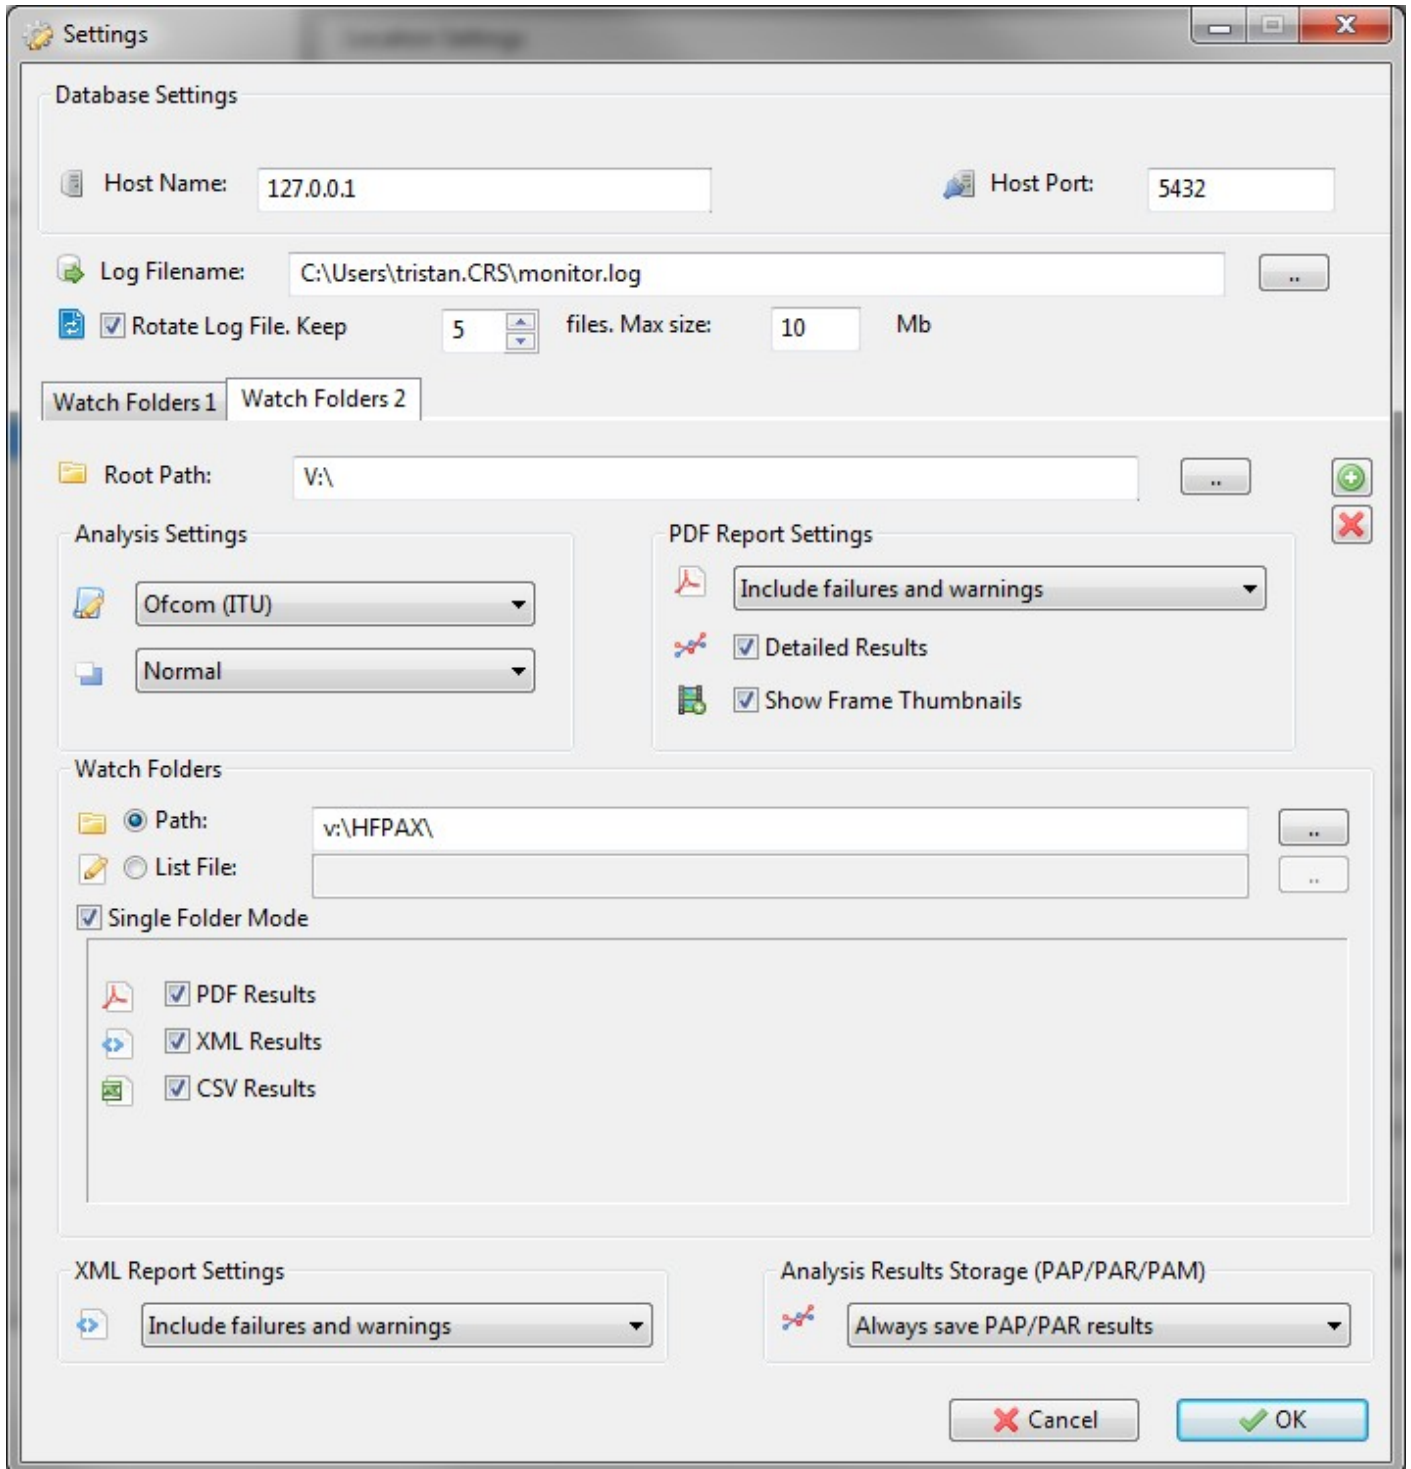
To enable this mode, select the *Single Folder Mode* checkbox.

Please note that whilst operating in this mode, the Job ID number is **not** prepended to the start of the filenames for the results or progress files.

Due to the fact that the movie assets will not be moved during and after analysis, the *Monitor* will not requeue a job in the single watch folder if at least one results file corresponding to the file in question is

present. If these results files are deleted or moved out of the single watch folder, then the file will be requeued.

The screenshot below shows an example settings screen for using a single watch folder.



Results Output Settings:

Below the watch folder settings there are some drop down boxes to select the criteria for generating the different results files, as well as options for inclusion of detailed results and/or thumbnails in the PDF certificate:

Advanced XML Results Configuration:

Should you require that a specific string be appended to the end of the filename for the XML results files, you can configure this by adding a string value to the Windows registry / OS X ini file as follows:

Windows:

Add a string key XML_Append :

```
HKEY_LOCAL_MACHINE\SOFTWARE\Cambridge Research Systems\HFPAXAnalyser\XML_
Append
```

OS X:

In the file /Library/HfpaPreferences/com.crsLtd.hfpaxmonitor.ini

Add a string to the[SETTINGS] section:

```
XML_Append = _results
```

The string contained for this key/value is then appended to the XML results filename. For example if "_results" is placed into the string, the XML filename would appear as below:

```
MovieName_results.xml
```

XML Token File

The XML Token File is an advanced way of submitting jobs to the HardingFPA-X System. The parameters for the job to be submitted are entered into an XML "token file", which is then copied to one of the *input* watch folders. The *Monitor* application processes the token and subsequently sets up a job according to its contents and submits it to the *Job Queue*.

The format of this file and an example is shown below (with example values):

```
<?xml version="1.0" encoding="ISO-8859-1" ?>
```

```

<hfpaxjob>

  <priority>0</priority>

  <file title="test" subtitle="subtitle" extra="extra"
  assetnumber="as12345">/Users/tristan/Movies/test.mov</file>

  <locations>

    <output jobid="1">/Users/shared/Output/</output>

    <pappar jobid="1"
    stor-
    age="always">/Users/shared/Results/PAPPAR</pappar>

    <pdf jobid="1" content="all" detailed="1"
    thumbnails="1">/Users/shared/Results/PDF</pdf>

    <csv jobid="1"
    storage="always">/Users/shared/Results/CSV</csv>

    <xml jobid="1"
    content="all">/Users/shared/Results/XML</xml>

  </locations>

</hfpaxjob>

```

In this file the job's priority value is given as the `priority` field and can take a value of between -10 and 10, where 0 is standard priority and is also the default value if no priority is given. Values greater than 10 or less than -10 will be capped to 10 or -10.

The movie file location is specified in the `<file>` tag (`/Users/tristan/Movies/test.mov`), with optional Material Description specified in the **title**, **subtitle**, **extra** and **assetnumber** fields. The Asset Number allows you to better integrate with automation systems by providing a reference to the movie that you may track. The Asset Number is a string (i.e. can contain letters as well as numbers) and is written to XML Progress files, XML Results files and also the PDF certificate Results files. All other Material Description fields will be added into the PAR and PDF results files that are created.

Next is the locations section, which describes the destinations of the results files. These are all optional, and results of the corresponding type will only be output if there is a path present in the relevant section. The first field in each of these, `jobid`, should be set to "1" if you require the Job ID number of the job to be prepended to the results file, e.g. `00059test.pdf`, and to "0" if this is not required, e.g. `test.pdf`.

The results directories are described by the following sections:

- `<output>` Move the movie file into the following folder. If you would like to set separate passed, failed and error folders, you may instead include the two tags `output_passed`, `output_failed` and `output_error` which must appear in this order:
 - `<output_passed jobid="1">/Users/shared/Passed/</output_passed>`
 - `<output_failed jobid="1">/Users/shared/Failed/</output_failed>`
 - `<output_error jobid="1">/Users/shared/Error/</output_error>`
- `<csv>` Create CSV Results in the following folder.
 - In this option, "storage" can be set to specify when the CSV results are to be saved: "always" to always save, "warning" to only save when a warning or failure occurs, "failure" to only save when the material fails, and "never" to never save CSV results.
- `<pappar>` Create PAP/PAR/PAM Results in the following folder.
 - In this option, "storage" can be set to specify when the PAP/PAR/PAM results are to be saved: "always" to always save, "warning" to only save when a warning or failure occurs, "failure" to only save when the material fails, and "never" to never save PAP/PAR/PAM results.
- `<pdf>` Create PDF Results in the following folder.
 - In this option, "content" can be set to specify which details are saved to the PDF certificate produced for this job: "all" to save warnings and failure information, or "failures" to only save failure information.
 - To show details of all failed frames in the PDF certificate, set "detailed" to 1, otherwise set to 0 (where the start and end frames of a warning/fail sequence are reported).
 - To show thumbnail images in the PDF certificate, set "thumbnails" to 1, otherwise set to 0.
- `<xml>` Create XML Results in the following folder.
 - In this option, "content" can be set to specify what details are saved to the XML results file produced for this job: "all" to save warnings and failure information, or "failures" to only save failure information.


HardingFPA-X Supervisor

This is a tool for the exclusive use of the system administrator, and can be run from as many computers as it is installed on. The *Supervisor* application allows administrative users to licence client computers to use the *Submit* as well as manage jobs that are currently in the job queue, or currently being analysed.

With this application, the administrator can also modify the priorities of jobs in the Job Queue and abort them.

On start up, the Supervisor Password must be entered to gain control of the database. The default password here is “**adminhfpax**”.

Please note that this password is specific to the database installation. It is recommended to change the Supervisor Password as soon as possible by selecting *Tools* → *Change Supervisor Password*.

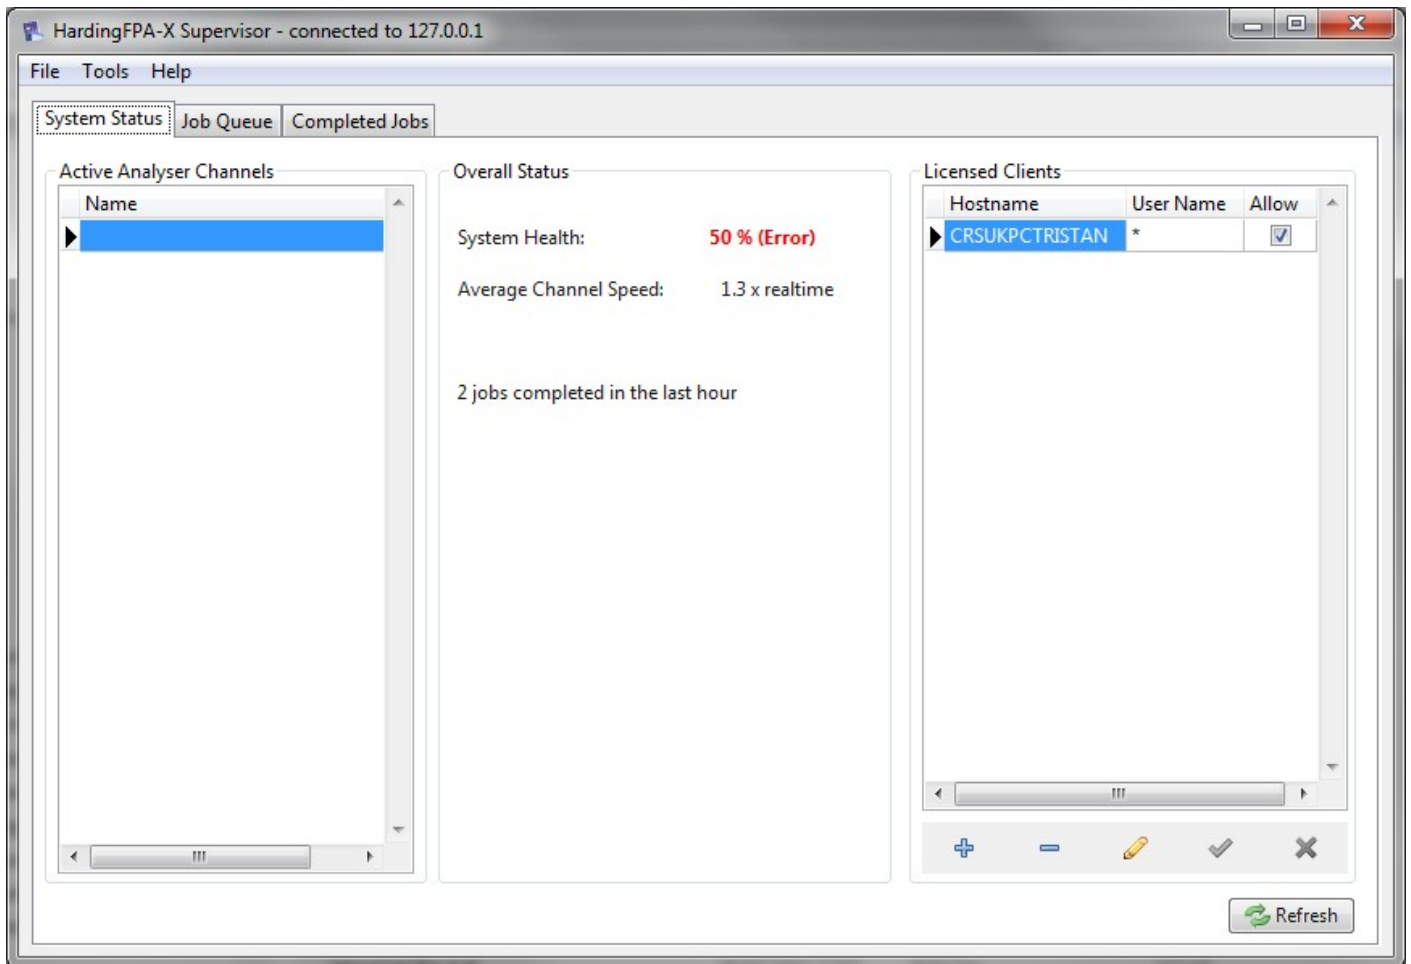


On the HardingFPA-X Supervisor interface, there are three tabs. The first of these is a general overview of the *System Status*, and shows various pieces of information about the system as a whole and allows the supervisor to be able to see immediately if there are any serious errors occurring in the system as well as viewing a list of all of the currently active *Analyser* channels and licensed client computers.

The *System Health* value is a percentage that describes the proportion of the last 20 jobs that completed successfully. Note that this does not refer to the pass/fail status of the jobs, but whether or not a serious error occurred meaning that the job could not be analysed at all.

The *Average Channel Speed* value is the average rate at which an *Analyser* channel will analyse jobs, so that you may see the expected speed of jobs being analysed at the moment. This is per channel, so if your system has an HardingFPA-X Analyser with two channels (as in a typical installation), the actual system throughput will be twice this value.

At the bottom of the *Overall Status* section, the number of jobs that have been completed in the last hour is shown.



The *Supervisor* also features a Job Queue, which shows all of the jobs that are waiting to be analysed or are currently being analysed – in the order in which they will be opened by the *Analysers*. The final tab (Completed Jobs) shows all of the jobs which have been analysed (or have been aborted or terminated with an error condition). Unlike the *Monitor* and *Submit*, this window shows all completed jobs, not just those from the last 14 days, and shows which user submitted which jobs and which *Analyser* analysed the job.

Licensing Clients

The *Supervisor* is able to display a list of all the licensed computers/users that are allowed to use the *Submit* application.

They are listed in the table on the right hand side of the first tab *System Status*, as shown below:

Licensed Clients

Hostname	User Name	Allow
CRSUKPCTRISTAN	*	<input checked="" type="checkbox"/>









The cells in this table can be edited to modify the licensed clients.

The database supports the * character for wildcards, so if for example "*" is entered for User Name, then all users at that computer will be licensed.

A "*" in both columns means licence all users on all computers.

To add a new user/computer combination, click the plus button () and edit the table, clicking the tick button () when done.

Job Priorities

By default, all jobs submitted to the system, whether via *Submit* or *Monitor* (watch folders etc), have an initial Job Priority of **0**. In effect this means the job is neither low priority nor high priority. A job's priority value can be between -10 and 10.

If all jobs have the same priority, then they will be analysed in order of submission, i.e. the lowest Job ID number first (please note that some later jobs may be analysed before earlier ones if the *Analyser(s)* configured for that job length is busy). If the queue contains jobs of mixed priorities, then they are analysed first in order of priority, then in order of job ID. For example if two jobs have priority **1** but all the rest have a priority of **0**, the priority **1** jobs will first be analysed in the usual order of their Job ID, after which all jobs of priority 0 will be analysed, also in order of Job ID, etc.

Using the Promote Job and Demote Job buttons under the Job Queue tab, the administrator may increase or decrease the priority of the selected job respectively. The change in priority is shown in the Job Priority column. Be aware that you cannot increase the priority to greater than 10, or decrease it to lower than -10. The change in the job order is reflected immediately in the order displayed in the Job Queue.

Promoting jobs is a useful tool for administrators to be able to advance emergency and/or time critical jobs to the top of the queue. Be aware that if a job is demoted to have a negative priority, it may not be analysed at all if a continual stream of jobs is submitted, as the newly submitted jobs will, by default, have a higher priority. To be sure, always set all job priorities back to 0 after manipulating the job queue for emergency jobs etc.

HardingFPA-X Submit

This application must be licensed by the administrator user of the HardingFPA-X system. This can be done in one of two ways:

1. Using the [Licensing Clients](#) panel in the *Supervisor* application to add the computer name to the list of licensed clients.
2. Running the Submit on an unlicensed computer, and logging in with the administrator's password (default "*adminhfpax*") to register it in the database.

This application is designed to be run on any number of licensed client computers and allows the user of that computer to submit jobs to the database ready for the *Analyser(s)* to process them. Once a job is completed, the Submit application allows the user to right-click on the completed job and launch the Viewer or a third party PDF reader to view the results of the analysis.

The Job Queue lists all of the currently pending jobs for all connected *Analysers* in the system, and the Completed Jobs list shows all of the completed jobs for the last 14 days.

The *Submit* application must be configured to point at the current root path in order to submit jobs and view results. Detailed instructions can be found in the HardingFPA-X System User's Manual.

If you would like to lock the settings for the *Submit* so that client users cannot modify them, you may set up the *Submit* how you wish and then set the preferences ini files or registry keys to read only for the users in question. On Windows Vista/Windows 7 this is provided automatically by the *Run As Administrator* option.

The database settings are stored in:

OS X: /Library/HfpaPreferences/com.crsLtd.hfpadatabase.ini

Windows: On Windows, the settings are stored in the registry key [HKEY_LOCAL_MACHINE\SOFTWARE\Cambridge Research Systems\HFPAXDatabase]. A restricted user will not be able to write to these unless the application is run as an administrator.

The database settings are shared between all applications installed on the same computer.

The general *Submit* settings are stored in:

OS X: /Library/HfpaPreferences/com.crsLtd.hfpasubmit.ini

Windows: On Windows, the settings are stored in the registry key [HKEY_LOCAL_MACHINE\SOFTWARE\Cambridge Research Systems\HFPAXSubmit]. A restricted user will not be able to write to these unless the application is run as an administrator.

HardingFPA Viewer

This is the part of the HardingFPA-X system which allows the client computers to view the results of completed analysis, displayed in the same way as in the standard version of the HardingFPA, with the additional benefit of being able to replay the image sections, and view Frame Mask data.

The *Viewer* is usually installed on the same computers as the *Submit* application and is launched automatically to view results.

Detailed instructions for the usage of HardingFPA-Viewer can be found in the HardingFPA-X System User's Manual.

Example Advanced Configurations

The HardingFPA-X System is very versatile, and can be set up in many configurations. For example it is entirely possible to have all the required Server, Supervisor and Client components on a single computer, but the flexibility provided by the HardingFPA-X System allows for more complex and useful configurations such as those detailed below. Please note that additional licences will need to be purchased in order to run multiple Analysers.

- **Typical Usage:** Running the main database, *Launcher*, *Monitor*, and *Analyser* on one powerful multi-core computer, and connecting multiple clients each with a *Submit* and *Viewer*. This configuration separates the tasks of analysis from the submission of jobs and viewing of results, leaving the main server computer with plenty of time to process jobs. Jobs can also be submitted via watch folders and XML tokens due to the presence of the *Monitor*.
- **Multiple Analyser Nodes:** Having a single computer running the database and *Monitor*, then having for example three computers all running separate *Analyser* instances (with *Launchers*). Multiple clients connect as before but this time the three dedicated *Analysers* together can process approximately three times as quickly as the scenario above.
- **Multiple Databases:** In circumstances whereby you need to have a number of distinct different job queues, you could have the database installed multiple times, once for each queue, and then have a set of applications connected to each.
- **Multiple Analyser and Monitor Applications:** It is possible to run the *Monitor* on several computers, and "link" the watch folders together, such that the output of one watch folder structure is the input of another watch folder structure. This can then be used in conjunction with more than one *Analyser*, each set to a different analysis standard, for instance.

Troubleshooting

Symptom	Solution
When running the <i>Analyser</i> on OS X, the error message "HASP SRM Protection System Feature not found (31)" appears, even though the HASP key is inserted.	The HASP key you are using is an older one, which requires a firmware update. Please see the HASP Firmware Update folder on the installation media provided with the HardingFPA-X.
Cannot log in to the HardingFPA-X <i>Supervisor</i> application.	The default password is "adminhfpax"
Forgot supervisor password for HardingFPA-X Supervisor.	Start the <i>Supervisor</i> application on the same computer as the database is installed on, and enter no password. The database will authenticate and you can change the password with Tools -> Change Supervisor password.
Job does not get analysed, but stays at the bottom of the Job Queue.	Check the priority of the job; it may have an unusually low priority. Promote the job to a higher priority using the Supervisor.
Jobs are not being analysed because the <i>Analyser</i> cannot find the movie files.	Make sure the <i>Analyser</i> has access to the files in question, and that the <i>Submit</i> and <i>Analyser</i> applications have the correct Root Path settings.
<i>Analysers</i> do not open files dropped into watch folders.	Check that the file extension is one of the supported extensions and that the case of all of the letters in the extension are the same (e.g. MOV, not MoV). Files dropped into watch folders must also not be read-only. Files must be larger than 2kb.
<i>Analysers</i> will not run, and give an error message such as "This version is older than the newest version in the system".	The system has been upgraded, but the <i>Analyser</i> in question has not. Upgrade the <i>Analyser</i> to the same version as the rest of the system.
The paths and database settings are correct, but the <i>Analyser(s)</i> will not analyse a particular job.	Does the computer analysing the job have the required codec? – check if the file can be viewed correctly in QuickTime. Also check if the frame size and frame rate are within the accepted range.
The Server installer on OS X hangs and does not	If you are upgrading, make sure all HardingFPA-X

complete.

applications on your entire network are not running.
This includes *Analysers*, *Monitors*, *Submits* and
Supervisors.

Licence Terms

This HardingFPA product utilises icons from the Silk icon set at <http://www.famfamfam.com/lab/icons/silk/> and the Fat Cow icon set at <http://www.fatcow.com/free-icons>.

Some plugins use external libraries and makes extensive use of the following persons' or companies' code:

FFmpeg - Copyright (c) 2000-2011 Fabrice Bellard, et al.

<http://www.ffmpeg.org/>

The source and build instructions for the included FFmpeg libraries are contained on the HardingFPA-X installation media.

FFmpeg is licensed under the LGPL Licence, duplicated below:

GNU LESSER GENERAL PUBLIC LICENSE

Version 2.1, February 1999

Copyright (C) 1991, 1999 Free Software Foundation, Inc.

51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

[This is the first released version of the Lesser GPL. It also counts as the successor of the GNU Library Public License, version 2, hence the version number 2.1.]

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public Licenses are intended to guarantee your freedom to share and change

free software--to make sure the software is free for all its users.

This license, the Lesser General Public License, applies to some specially designated software packages--typically libraries--of the Free Software Foundation and other authors who decide to use it. You can use it too, but we suggest you first think carefully about whether this license or the ordinary General Public License is the better strategy to use in any particular case, based on the explanations below.

When we speak of free software, we are referring to freedom of use, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish); that you receive source code or can get it if you want it; that you can change the software and use pieces of it in new free programs; and that you are informed that you can do these things.

To protect your rights, we need to make restrictions that forbid distributors to deny you these rights or to ask you to surrender these rights. These restrictions translate to certain responsibilities for you if you distribute copies of the library or if you modify it.

For example, if you distribute copies of the library, whether gratis or for a fee, you must give the recipients all the rights that we gave you. You must make sure that they, too, receive or can get the source code. If you link other code with the library, you must provide complete object files to the recipients, so that they can relink them with the library after making changes to the library and recompiling it. And you must show them these terms so they know their rights.

We protect your rights with a two-step method: (1) we copyright the library, and (2) we offer you this license, which gives you legal permission to copy, distribute and/or modify the library.

To protect each distributor, we want to make it very clear that there is no warranty for the free library. Also, if the library is modified by someone else and passed on, the recipients should know that what they have is not the original version, so that the original author's reputation will not be affected by problems that might be introduced by others.

Finally, software patents pose a constant threat to the existence of any free program. We wish to make sure that a company cannot effectively restrict the users of a free program by obtaining a restrictive license from a patent holder. Therefore, we insist that any patent license obtained for a version of the library must be consistent with the full freedom of use specified in this license.

Most GNU software, including some libraries, is covered by the ordinary GNU General Public License. This license, the GNU Lesser General Public License, applies to certain designated libraries, and is quite different from the ordinary General Public License. We use this license for certain libraries in order to permit linking those libraries into non-free programs.

When a program is linked with a library, whether statically or using a shared library, the combination of the two is legally speaking a combined work, a derivative of the original library. The ordinary General Public License therefore permits such linking only if the

entire combination fits its criteria of freedom. The Lesser General Public License permits more lax criteria for linking other code with the library.

We call this license the "Lesser" General Public License because it does Less to protect the user's freedom than the ordinary General Public License. It also provides other free software developers Less of an advantage over competing non-free programs. These disadvantages are the reason we use the ordinary General Public License for many libraries. However, the Lesser license provides advantages in certain special circumstances.

For example, on rare occasions, there may be a special need to encourage the widest possible use of a certain library, so that it becomes a de-facto standard. To achieve this, non-free programs must be allowed to use the library. A more frequent case is that a free library does the same job as widely used non-free libraries. In this case, there is little to gain by limiting the free library to free software only, so we use the Lesser General Public License.

In other cases, permission to use a particular library in non-free programs enables a greater number of people to use a large body of free software. For example, permission to use the GNU C Library in non-free programs enables many more people to use the whole GNU operating system, as well as its variant, the GNU/Linux operating system.

Although the Lesser General Public License is Less protective of the users' freedom, it does ensure that the user of a program that is

linked with the Library has the freedom and the wherewithal to run that program using a modified version of the Library.

The precise terms and conditions for copying, distribution and modification follow. Pay close attention to the difference between a "work based on the library" and a "work that uses the library". The former contains code derived from the library, whereas the latter must be combined with the library in order to run.

GNU LESSER GENERAL PUBLIC LICENSE

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License Agreement applies to any software library or other program which contains a notice placed by the copyright holder or other authorized party saying it may be distributed under the terms of this Lesser General Public License (also called "this License").

Each licensee is addressed as "you".

A "library" means a collection of software functions and/or data prepared so as to be conveniently linked with application programs (which use some of those functions and data) to form executables.

The "Library", below, refers to any such software library or work which has been distributed under these terms. A "work based on the Library" means either the Library or any derivative work under copyright law: that is to say, a work containing the Library or a portion of it, either verbatim or with modifications and/or translated straightforwardly into another language. (Hereinafter, translation is included without limitation in the term "modification".)

"Source code" for a work means the preferred form of the work for

making modifications to it. For a library, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the library.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running a program using the Library is not restricted, and output from such a program is covered only if its contents constitute a work based on the Library (independent of the use of the Library in a tool for writing it). Whether that is true depends on what the Library does and what the program that uses the Library does.

1. You may copy and distribute verbatim copies of the Library's complete source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and distribute a copy of this License along with the Library.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Library or any portion of it, thus forming a work based on the Library, and copy and distribute such modifications or work under the terms of Section 1

above, provided that you also meet all of these conditions:

- a) The modified work must itself be a software library.
- b) You must cause the files modified to carry prominent notices stating that you changed the files and the date of any change.
- c) You must cause the whole of the work to be licensed at no charge to all third parties under the terms of this License.
- d) If a facility in the modified Library refers to a function or a table of data to be supplied by an application program that uses the facility, other than as an argument passed when the facility is invoked, then you must make a good faith effort to ensure that, in the event an application does not supply such function or table, the facility still operates, and performs whatever part of its purpose remains meaningful.

(For example, a function in a library to compute square roots has a purpose that is entirely well-defined independent of the application. Therefore, Subsection 2d requires that any application-supplied function or table used by this function must be optional: if the application does not supply it, the square root function must still compute square roots.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Library, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based

on the Library, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Library.

In addition, mere aggregation of another work not based on the Library with the Library (or with a work based on the Library) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may opt to apply the terms of the ordinary GNU General Public License instead of this License to a given copy of the Library. To do this, you must alter all the notices that refer to this License, so that they refer to the ordinary GNU General Public License, version 2, instead of to this License. (If a newer version than version 2 of the ordinary GNU General Public License has appeared, then you can specify that version instead if you wish.) Do not make any other change in these notices.

Once this change is made in a given copy, it is irreversible for that copy, so the ordinary GNU General Public License applies to all subsequent copies and derivative works made from that copy.

This option is useful when you wish to copy part of the code of the Library into a program that is not a library.

4. You may copy and distribute the Library (or a portion or derivative of it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange.

If distribution of object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place satisfies the requirement to distribute the source code, even though third parties are not compelled to copy the source along with the object code.

5. A program that contains no derivative of any portion of the Library, but is designed to work with the Library by being compiled or linked with it, is called a "work that uses the Library". Such a work, in isolation, is not a derivative work of the Library, and therefore falls outside the scope of this License.

However, linking a "work that uses the Library" with the Library creates an executable that is a derivative of the Library (because it contains portions of the Library), rather than a "work that uses the library". The executable is therefore covered by this License.

Section 6 states terms for distribution of such executables.

When a "work that uses the Library" uses material from a header file that is part of the Library, the object code for the work may be a derivative work of the Library even though the source code is not.

Whether this is true is especially significant if the work can be

linked without the Library, or if the work is itself a library. The threshold for this to be true is not precisely defined by law.

If such an object file uses only numerical parameters, data structure layouts and accessors, and small macros and small inline functions (ten lines or less in length), then the use of the object file is unrestricted, regardless of whether it is legally a derivative work. (Executables containing this object code plus portions of the Library will still fall under Section 6.)

Otherwise, if the work is a derivative of the Library, you may distribute the object code for the work under the terms of Section 6. Any executables containing that work also fall under Section 6, whether or not they are linked directly with the Library itself.

6. As an exception to the Sections above, you may also combine or link a "work that uses the Library" with the Library to produce a work containing portions of the Library, and distribute that work under terms of your choice, provided that the terms permit modification of the work for the customer's own use and reverse engineering for debugging such modifications.

You must give prominent notice with each copy of the work that the Library is used in it and that the Library and its use are covered by this License. You must supply a copy of this License. If the work during execution displays copyright notices, you must include the copyright notice for the Library among them, as well as a reference directing the user to the copy of this License. Also, you must do one of these things:

- a) Accompany the work with the complete corresponding machine-readable source code for the Library including whatever changes were used in the work (which must be distributed under Sections 1 and 2 above); and, if the work is an executable linked with the Library, with the complete machine-readable "work that uses the Library", as object code and/or source code, so that the user can modify the Library and then relink to produce a modified executable containing the modified Library. (It is understood that the user who changes the contents of definitions files in the Library will not necessarily be able to recompile the application to use the modified definitions.)
- b) Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that (1) uses at run time a copy of the library already present on the user's computer system, rather than copying library functions into the executable, and (2) will operate properly with a modified version of the library, if the user installs one, as long as the modified version is interface-compatible with the version that the work was made with.
- c) Accompany the work with a written offer, valid for at least three years, to give the same user the materials specified in Subsection 6a, above, for a charge no more than the cost of performing this distribution.
- d) If distribution of the work is made by offering access to copy from a designated place, offer equivalent access to copy the above specified materials from the same place.

e) Verify that the user has already received a copy of these materials or that you have already sent this user a copy.

For an executable, the required form of the "work that uses the Library" must include any data and utility programs needed for reproducing the executable from it. However, as a special exception, the materials to be distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

It may happen that this requirement contradicts the license restrictions of other proprietary libraries that do not normally accompany the operating system. Such a contradiction means you cannot use both them and the Library together in an executable that you distribute.

7. You may place library facilities that are a work based on the Library side-by-side in a single library together with other library facilities not covered by this License, and distribute such a combined library, provided that the separate distribution of the work based on the Library and of the other library facilities is otherwise permitted, and provided that you do these two things:

a) Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities. This must be distributed under the terms of the Sections above.

b) Give prominent notice with the combined library of the fact that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.

8. You may not copy, modify, sublicense, link with, or distribute the Library except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, link with, or distribute the Library is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

9. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Library or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Library (or any work based on the Library), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Library or works based on it.

10. Each time you redistribute the Library (or any work based on the Library), the recipient automatically receives a license from the original licensor to copy, distribute, link with or modify the Library subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein.

You are not responsible for enforcing compliance by third parties with this License.

11. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Library at all. For example, if a patent license would not permit royalty-free redistribution of the Library by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply, and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

12. If the distribution and/or use of the Library is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Library under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

13. The Free Software Foundation may publish revised and/or new versions of the Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Library specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Library does not specify a license version number, you may choose any version ever published by the Free Software Foundation.

14. If you wish to incorporate parts of the Library into other free programs whose distribution conditions are incompatible with these, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our

decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE LIBRARY IS WITH YOU. SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Libraries

If you develop a new library, and you want it to be of the greatest possible use to the public, we recommend making it free software that everyone can redistribute and change. You can do so by permitting redistribution under these terms (or, alternatively, under the terms of the ordinary General Public License).

To apply these terms, attach the following notices to the library. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

<one line to give the library's name and a brief idea of what it does.>

Copyright (C) <year> <name of author>

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2.1 of the License, or (at your option) any later version.

This library is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU Lesser General Public License for more details.

You should have received a copy of the GNU Lesser General Public License along with this library; if not, write to the Free Software Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA

Also add information on how to contact you by electronic and paper mail.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the library, if

necessary. Here is a sample; alter the names:

Yoyodyne, Inc., hereby disclaims all copyright interest in the

library `Frob' (a library for tweaking knobs) written by James Random Hacker.

<signature of Ty Coon>, 1 April 1990

Ty Coon, President of Vice

That's all there is to it!