

Quantum[®]

User's Guide

StorNext Timecode-Based Partial File Retrieval 1.1



Document Title, 6-67544-01 Rev A, January 2012, Product of USA.

Quantum Corporation provides this publication “as is” without warranty of any kind, either express or implied, including but not limited to the implied warranties of merchantability or fitness for a particular purpose. Quantum Corporation may revise this publication from time to time without notice.

COPYRIGHT STATEMENT

© 2012 Quantum Corporation. All rights reserved.

Your right to copy this manual is limited by copyright law. Making copies or adaptations without prior written authorization of Quantum Corporation is prohibited by law and constitutes a punishable violation of the law.


TRADEMARK STATEMENT

Quantum, the Quantum logo, DLT, DLTtape, the DLTtape logo, Scalar, and StorNext are registered trademarks of Quantum Corporation, registered in the U.S. and other countries.

Backup. Recovery. Archive. It’s What We Do., the DLT logo, DLTSage, DXi, DXi-Series, Dynamic Powerdown, FastSense, FlexLink, GoVault, MediaShield, Optyon, Pocket-sized. Well-armored, SDLT, SiteCare, SmartVerify, StorageCare, Super DLTtape, SuperLoader, and Vision are trademarks of Quantum.

LTO and Ultrium are trademarks of HP, IBM, and Quantum in the U.S. and other countries. All other trademarks are the property of their respective companies.

Specifications are subject to change without notice.



Contents

Chapter 1	Introduction to Timecode-Based Partial File Retrieval	1
<hr/>		
Chapter 2	How Timecode-Based Partial File Retrieval Works	3
	Partial File Retrieval Indexer	4
	Partial File Retrieval Service	5
	Folder Structure	5
	Issuing a Request	5
	Monitoring Requests.	6
	PFR Workflow	7
	Codecs.	11
	File Format Specifications	13
	QuickTime Specification	13
	MXF Specification	13
	Minimum Restore Sizes.	13
<hr/>		
Chapter 3	Installation	15
	Installation Components and Install Options.	15


Upgrading from an Earlier Version	16
Uninstalling a Previous Version	16
Installing v1.1	17
Entering Web Services Information	17
Prerequisites	17
Supported Platforms	17
User Account Control in Windows 7 and Windows Server 2008	18
StorNext	18
QuickTime (is not required)	19
Bonjour	19
User Accounts	20
.NET	21
Running the Installer Program	24
Installation Type Options	24
Licensing your Product	26
Overview	26
Activating the Initial Temporary License	26
Obtaining a Permanent License	27
Installing the Permanent License	27
Extending a Temporary License	28
Migrating a License to a New Machine	29

Chapter 4

Configuration	31
Virtual Machine From Earlier Version of Quantum PFR	31
Primary Configuration	31
Welcome	32
Archive	32
Job Results Folder	33
Work Folders	34
System Configuration	35
Indexer	37
Summary	38
Complete	38
Indexer Configuration and Monitoring	39
Configuration	40
Monitoring	41
Making Changes in a PFR Active/Active Cluster Configuration	42
Starting / Stopping the Synchronization Service	43

Chapter 5	PFR Active/Active Clustering	45
	Overview	45
	PFR Server Redundancy	46
	Hardware Configurations	49
	Location of PFR Indexers	51
	Resubmitting Requests After Control Server Failover	51
Chapter 6	Status and GUI Application	55
	Anatomy of Status GUI	55
	Manipulating the Retrievals	56
	PFR Log Viewer	57
Chapter 7	Request EDL Format	61
	Sample	61
	Notes	62
	Special Characters	62
Chapter 8	Performance and Scalability	65
	Performance	65
	Indexer	65
	Note on Existing Archives	66
	Partial File Retrieval via Index from Tape	66
	Very Short Partial Retrievals	66
	Longer Partial Retrievals	67
	File and Folder Structure Limitations	67
	Quantum PFR Indexer	67
	Quantum PFR Server	68
	Scalability	68
	Distributing the Load	68
	Dedicated Tape Drives	68
	Suggested Platform	69

Chapter 9	Index Web Service API	71
	API Detail	72
	Indexing a File.	72
	Requesting File Index Status	73
	Returned XML.	74



Figures



Figure 1	PFR Steps 1 - 2	7
Figure 2	PFR Steps 3 - 5	8
Figure 3	PFR Steps 6 - 7	9
Figure 4	PFR Steps 8 - 9	10
Figure 5	Server Manager	21
Figure 6	Add Features	22
Figure 7	Install .NET Framework	23
Figure 8	PFR Installer	24
Figure 9	Welcome Screen	32
Figure 10	Archive Screen	33
Figure 11	Job Results Screen	34
Figure 12	Work Folders Screen	35
Figure 13	Transfer Engines Screen	36
Figure 14	PFR Indexer Screen	37
Figure 15	Summary Screen	38
Figure 16	Complete Screen	39
Figure 17	PFR Index Configuration	40
Figure 18	Indexer Log Window	42
Figure 19	PFR Server Redundancy	46
Figure 20	Combined Control Server & Engine	49
Figure 21	Combined Control Server & Engine with Additional Engine	50

Figures

Figure 22	Redundant Control Server with 2+ Engines	50
Figure 23	Status GUI Screen	56
Figure 24	Blank Log Viewer Screen.	57
Figure 25	Populated Log Viewer Screen	58
Figure 26	Merged Log Viewer Screen	59



Chapter 1

Introduction to Timecode-Based Partial File Retrieval

Quantum Timecode-Based Partial File Retrieval (PFR) allows you to quickly retrieve and utilize segments of large media files—rather than the entire file—based on timecode parameters.

Segments from large media files that have been archived to tape using Quantum StorNext Storage Manager can be retrieved using a simple XML request file. Quantum Partial File Retrieval instructs StorNext to retrieve only the necessary portion from tape, and then presents that segment as a complete MXF or QuickTime file.

If the file has not yet been truncated by StorNext Storage Manager, the result will still be the same. However, in this case the Partial File will be generated directly from the complete media file on disk.

The Quantum Partial File Retrieval feature requires a dedicated Windows-based server on which to run. For larger scale installations, the Quantum Partial File Retrieval functionality can be distributed across multiple Windows server machines.



Chapter 2

How Timecode-Based Partial File Retrieval Works

Quantum Timecode-Based Partial File Retrieval consists of two primary components: an indexer which indexes media files before they are archived to tape and truncated, and a Partial File Retrieval service that performs the requested retrievals.

Partial File Retrieval is designed to work exclusively with QuickTime and MXF video files, which makes PFR suitable for applications which reference video by timecode.

The QuickTime or MXF wrapper for a retrieved file will be re-constructed during the Partial File Retrieval process, and as such retrieved files will not be byte-for-byte replicas of the original. However, the video frames and audio samples in the retrieved files are not transcoded and therefore there is no generation loss from the retrieval process. The only exception to this is for retrieving MPEG Long GOP video (for example, Sony XDCAM HD 50), where a small number of frames are transcoded at the start and end of the retrieval in order to reconstruct a coherent GOP structure.

Partial File Retrieval Indexer

The Partial File Retrieval Indexer looks for new media. When it finds a new media file it generates an index for it, in effect mapping timecode values in the original media file to byte offsets. This allows subsequent timecode-based requests to retrieve only the necessary portions of a large media file from tape.

New in Quantum PFR v1.1.0 is the ability to create indexes on demand (via a Web Service API) rather than automatically. This ability particularly suits environments where the archival process is under the control of an Asset Management system.

The index files, which are written as QuickTime reference files, are generated in a parallel folder structure on the managed storage. The two file locations below represent a MXF based media file and its associated index file.

```
\\StorNext\snfs1\Media\Project\Movie.mxf
```

```
\\StorNext\snfs1\PFR-INDEX\Media\Project\Movie.mov
```

Note: The index files do not contain any media file content, and are therefore very small in comparison with the files they index.

It is important that the indexed files themselves are never truncated. The only files that should be set for truncation are source media files. These media files should be set to truncate no sooner than 24 hours.

The Indexer does not attempt to index files currently being written. Therefore, the first time the indexer encounters the file, there may be some delay before a particular file is indexed while the indexer is trying to establish if the file is currently being written to. The delay could be significant if explicit indexing is not chosen, measured in hours or longer if many large files must be indexed.

Partial File Retrieval Service

Folder Structure

Within the Quantum Partial File Retrieval Configuration Tool, it is possible to set up multiple sets of folders for different request types, such as News, Sports, Weather and so on. For each request type there would be four directories. For example, a set of folders for a News request type might look like this:

```
\\StorNext\snfs1\Requests\News
```

```
\\StorNext\snfs1\Requests\News\PFR-REQUEST
```

```
\\StorNext\snfs1\Requests\News\success
```

```
\\StorNext\snfs1\Requests\News\fail
```

The configuration tool creates these folders for you and allows you to set an output format: QuickTime Movie or MXF.

The four folders shown in the example are used in the following manner:

- 1 The parent folder (in this case *News*) is where the resultant file extract will be placed once retrieved and created;
- 2 The *PFR-REQUEST* folder is the location to place the XML-formatted request file;
- 3 Once the request has been completed, the XML request file will be moved to the *success* folder;
- 4 If an error occurred in retrieving the partial file, the XML request file will be moved to the *fail* folder.

Issuing a Request

A Partial File Retrieval request is issued in the form of a small XML file, the format of which can be seen in the example later in this User Guide. The XML file simply describes a movie (via a full UNC path) with an in and out timecode value.

Quantum PFR will handle a number of the requests simultaneously with the overall retrieval performance dependant upon the quantity of tape drives available.

The number of tape drives limits how many requests can be processed simultaneously. Once all tape drives are busy, additional requests will be queued until a tape drive becomes available.

Note: If the XML request file is missing certain required fields, a variety of error conditions may result, depending on which field is missing.

Examples of missing fields include but are not limited to: frame rate, NumberVideoTracks, UID, Title, File, Start, and End.

As the XML request file is typically produced by applications, not individuals, missing fields would not normally occur.

The actual error condition may vary depending on which field is missing.

Monitoring Requests


Requests will be processed automatically once the request file is detected. The Partial File Retrieval service automatically detects the XML request file, appends the processed request to the queue of pending requests, and processes them in the order received. XML request files are moved to the success or fail folder as appropriate based on the status of the request.

Use the PFR status GUI to monitor requests, cancel pending or active requests or change the priority of the pending requests.

PFR Workflow

This section describes the steps in the life cycle of a movie in a PFR environment. Circled numbers in the illustrations correspond to the step number.

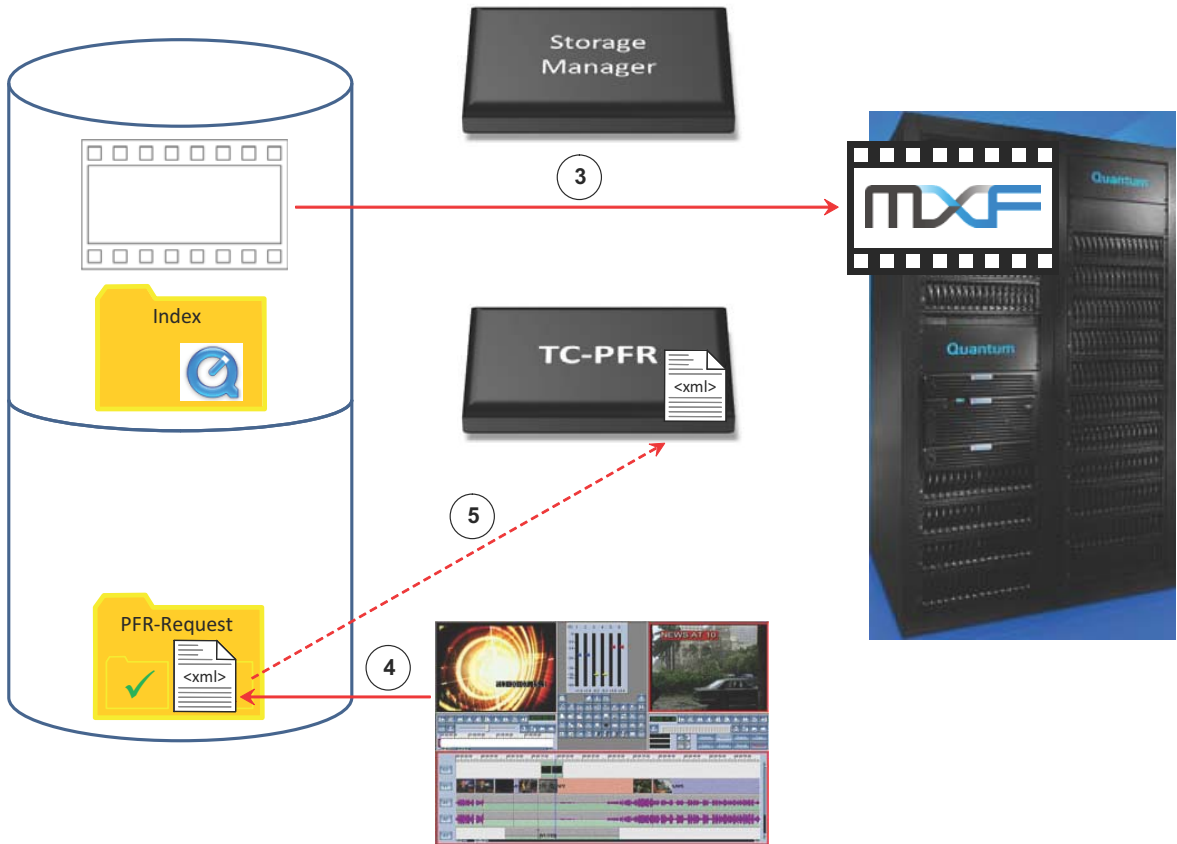
Figure 1 PFR Steps 1 - 2



Step 1: An MXF (or QuickTime) file is written into the managed file system.

Step 2: Frames are read from this new file and converted into a QuickTime reference index file.

Figure 2 PFR Steps 3 - 5




Step 3: Some time later, according to the policies that have been configured, the Storage Manager copies the file to the tape archive and truncates the file on disk.

Step 4: A third-party media asset management (MAM) system generates an XML request for a partial file and writes it into the request folder.

Step 5: PFR parses the file and converts it to a job retrieval request.


Figure 3 PFR Steps 6 - 7



Step 6: PFR locates the required frames in the index file and converts these to byte offset data for Storage Manager.

Step 7: Storage Manager performs a partial restore from tape, and writes a temporary file to the destination directory.

Figure 4 PFR Steps 8 - 9



Step 8: PFR then immediately rewraps this temporary file back to the requested MXF or QuickTime file format.

Step 9: The request XML is dropped into the success folder to indicate that a successful restore has occurred.

Codecs

Segments of large media files that have been archived to tape using StorNext Storage Manager can be requested using a simple XML request file. Timecode-Based Partial File Retrieval instructs StorNext to retrieve only the necessary portion from tape, and then presents that segment as a complete MXF or QuickTime file.

Following is a list of MXF and QuickTime file and media formats that have been tested with Partial File Retrieval:

Codec	NTSC	PAL	720p/50	720p/59.94	1080i/25	1080i/29.97
DV25	✓	✓	n/a	n/a	n/a	n/a
DVCPro	✓	✓	n/a	n/a	n/a	n/a
DVCPro50	✓	✓	n/a	n/a	n/a	n/a
DVCPro HD	n/a	n/a	✓	✓	✓	✓
IMX 30	✓	✓	n/a	n/a	n/a	n/a
IMX 40	✓	✓	n/a	n/a	n/a	n/a
IMX 50	✓	✓	n/a	n/a	n/a	n/a
MPEG2 I-Frame	n/a	n/a	✓	✓	✓	✓
MPEG2 Long GOP @ 25 Mbps	✓	✓	✓	✓	✓	✓
MPEG2 Long GOP @200 Mbps	✓	✓	✓	✓	✓	✓
ProRes 422 (Proxy) ProRes 422 (LT) ProRes 422 ProRes 422 (HQ) *	✓	✓	✓	✓	✓	✓
XDCAM 35	n/a	n/a	✓	✓	✓	✓

Codec	NTSC	PAL	720p/50	720p/59.94	1080i/25	1080i/29.97
XDCAM 50	n/a	n/a	✓	✓	✓	✓
DNxHD 120/145	n/a	n/a	MXF only	MXF only	MXF only	MXF only
DNxHD 185/220	n/a	n/a	MXF only	MXF only	MXF only	MXF only
DNxHD 185(x)/220(X)	n/a	n/a	MXF only	MXF only	MXF only	MXF only

* All ProRes codecs also support the following in addition to the values listed in the table: 1080p/23.976, 1080p/25 and 1080p/29.97.

Note: Retrieved 720p59.94 XDCAM HD 50 is not optimized for playback. For this one particular codec if you drag the restored PFR clip into a Final Cut Pro bin, you receive the “Not Optimized” message. However, after the warning is dismissed Final Cut Pro automatically optimizes the media.

When the “Not Optimized” message first appears, you have the option to either click OK, or to avoid future warnings (“Don’t warn me about this in the future.”) If you click OK, the media is automatically optimized and is playable in Final Cut Pro.

If you choose the avoid future warnings option and you later use the same codec again and drag a restored clip into the Final Cut Pro bin, the “Not Optimized” message is not displayed but the media is automatically optimized in the background and subsequently playable in Final Cut Pro. Quantum recommends choosing the option to never display future warnings when you first see the “Not Optimized” message.

File Format Specifications

This section contains information about QuickTime and MXF specifications as they pertain to the Partial File Retrieval feature.

QuickTime Specification

For ProRes, PFR 1.1 writes QuickTime-compliant files with QuickTime specification 2010-05-03.

Note that compliance alone does not guarantee playback in other devices. PFR 1.1 was evaluated with QuickTime Players version 7 and 10, and with Final Cut Pro.

MXF Specification

PFR 1.1 writes IMX according to SMPTE 386M D-10, and other MXF formats according to SMPTE RDD9.

Note: There may be differences between what the camera generates and what PFR 1.1 generates.

Minimum Restore Sizes

For fixed frame media such as IMX, DNxHD, DVCpro, etc., PFR can restore a single frame for either MXF or QuickTime wrappers.

For LGOP media such as XDCAMHD50, PFR can restore a single frame for MXF, but the minimum PFR can restore for a QuickTime wrapper is two frames.

Chapter 2: How Timecode-Based Partial File Retrieval Works File Format Specifications



Chapter 3 Installation

Installation Components and Install Options

Quantum Partial File Retrieval consists of the following software components:

- **PFR Control Server and Configuration GUI:** This is the core of the system and the tool that configures it
- **PFR Transfer Engine:** This is the engine that retrieves and creates the partial file
- **PFR Status GUI:** This is the user application which allows monitoring and modifying the queue of pending and in-progress retrievals
- **Index Engine and Configuration / Status GUI:** This is the service that indexes media files before they are truncated

A typical installation of Quantum Partial File Retrieval installs all of the above components on a single Windows-based server machine.

This is fine for most installations, but for larger installations the components can be distributed across several Windows servers. For more information, see [Chapter 5, PFR Active/Active Clustering](#).

Upgrading from an Earlier Version

Uninstalling a Previous Version

Prior to installing Quantum Partial File Retrieval v1.1, any prior version must be uninstalled.

Note: Uninstalling PFR does not remove any of the index files that were generated with an earlier version, nor does it remove the existing PFR configuration.

On the PFR server machine, go to services and select the synchronization service, tab to the log on view and make a note of the User that this service is running as (you will need to know this user and its password to complete the re-installation later)

To uninstall, find the PFR Components to uninstall:

- On Windows Server 2008 R2, from the Control Panel select **Uninstall a Program**

If the original installation was a Full install, there will be only one item to uninstall, labeled "PFR". If the original installation was performed by installing the component parts, uninstall any instances of "PFR Client", "PFR Server", "PFR Engine" and "PFR Indexer".

Once the component(s) have been uninstalled, Quantum recommends rebooting the Server to remove any temporary files before installing v1.1. (See below).

Note: In previous PFR versions the Virtual Machine ran the StorNext API (SNAPI) interface. Virtual Machine is now superfluous, so after the PFR upgrade the VMWare software can be removed from the Server unless it is needed for some other purpose.

All other prerequisites for v1.1 are the same as those for previous versions. All component parts of Quantum PFR must be upgraded to the same version.

Installing v1.1

Installing v1.1 should be performed as detailed below. Once installation is complete, it is important to run through the PFR Configuration Tool to re-save the Configuration.

Be sure to check the **Maximum Simultaneous Transfers** value. This value may need to be re-entered with the newer version of the software.

Entering Web Services Information

If you are upgrading from PFR 1.0.x, PFR 1.1 may not function correctly if you do not update the information in the “StorNext Web Service” section of the Archive screen. (For more information about this section of the Archive screen, refer to [Archive](#) on page 32.)

To access the Archive screen during configuration, follow these steps:

- 1 Select **Archive** in the **Work Folders** screen.
- 2 Click **Edit**.
- 3 Enter the Web Service information.
- 4 Click **Next** to save your changes and proceed with configuration.

Prerequisites

Supported Platforms

At this time Timecode-based Partial File Retrieval is supported only on the following platforms.

The PFR server and transfer engine can be installed on the following platforms:

- Windows Server 2008 R2 64-bit

PFR clients (Status GUI) can be installed and run on the following platforms:

- Windows Server 2008 32-bit and 64-bit
- Windows XP
- Windows Vista
- Windows 7

User Account Control in Windows 7 and Windows Server 2008

User Account Control must be set to **Never Notify** or **Off**.

To open User Account Control Settings on in Windows 7 or Windows Server 2008, do the following:

- 1 From the Start menu, type **UAC** in the **Search programs and Files** box.
- 2 Set the slider to **Never Notify**.
- 3 Restart the machine.

Caution: Settings are implemented only on a physical restart of the Windows machine.

StorNext

The StorNext metadata controller must have Storage Manager installed. PFR also uses the StorNext Web Services API (WSAPI), which is automatically installed with StorNext.

The PFR server must have the StorNext client software installed, and the client must be at the same StorNext revision level as the metadata controller (MDC).

File systems containing PFR video data must be mounted Read/Write as drive letters on the PFR Server. For details about installing and configuring the Windows client, see the *StorNext Installation Guide*.

Sharing the StorNext File System

Each StorNext file system should be shared via SMB (standard Windows server message block protocol). The security attributes of the Windows share on the file system must be set to “Full Control.”

In Windows Explorer, right-click on each StorNext file system drive letter and share the drive. The permissions of each share must be explicitly set to “Full Control” since the default value is Read-Only.

Important: Configuration information is shared between the machines that make up the PFR system (Control Servers and Transfer Engines). It is therefore important that share names and drive mappings are the same on all machines that form part of the system.

Defining Your Web Services Password

Follow these steps to define your Web Services password:

- 1 SSH to your StorNext metadata controller (MDC).
- 2 Log in as root.
- 3 Create the file by typing `vi /usr/adic/.snapipassword`
The file will initially be blank.
- 4 Put your desired password into the file you just created.
- 5 Save the file by pressing `shift + zz`.

After saving the file it is not necessary to restart cvfs service.

QuickTime (is not required)

Quantum PFR will both read and write media files wrapped as QuickTime files. It is NOT necessary to install any Apple QuickTime components to enable all of the QuickTime functionality within Quantum PFR.

Bonjour

Bonjour, also known as zero-configuration networking, enables automatic discovery of computers, devices, and services on IP networks. Bonjour uses industry-standard IP protocols to allow devices to automatically discover each other without needing to enter IP addresses or configure DNS servers.

Components of Quantum Partial File Retrieval will use Bonjour, if it is installed, to make configuration of the Client software easier. It is therefore recommended (but not required) that Bonjour be installed prior to installing Quantum Partial File Retrieval components. You can find Bonjour here:

http://support.apple.com/downloads/Bonjour_for_Windows

Note: Bonjour should be installed on any machine running the PFR Control Server software as well as any machine running the PFR Status Client software.

If you do not install Bonjour, you must enter the unc name of the Control Servers in a PFR Active/Active Cluster environment in the PFR Configuration Control Servers section.

User Accounts

The PFR Control Server and the PFR Transfer Engine both require a specific user account to run the Services. This user account must have full read and write access to the StorNext storage in order to perform partial file retrievals. The same user account is used for both components.

If you are performing an installation where some components are installed onto separate server machines, the user account name must be the same on both the machine which has PFR Control Server and the machine onto which the PFR Transfer Engine is installed. Remember, this must be a local user account and not a domain user account.


Note: On Windows platforms your local user account **MUST** have Administrator privileges. PFR may not operate correctly if your user account is not recognized as a local administrator.

.NET

This installation step is necessary on Windows Server 2008. (This operating system comes with .NET installed but disabled.) Before beginning the PFR installation, you must enable the .NET feature.


- 1 Launch the Windows Server Manager.

Figure 5 Server Manager




2 Go to Features and then click **Add Features**.

Figure 6 Add Features



3 Select .NET Framework and then click Install.

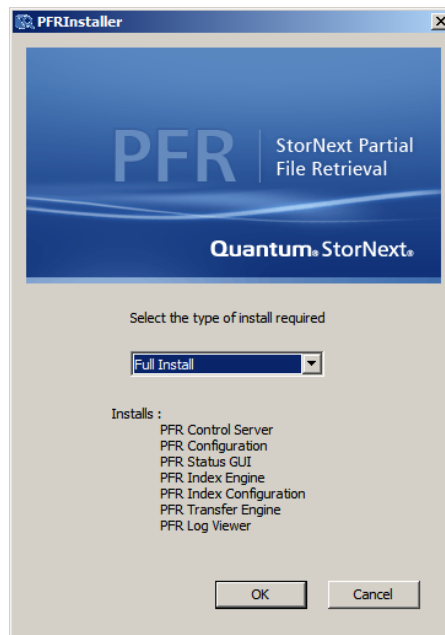
Figure 7 Install .NET Framework



Running the Installer Program

The installation starts with an options screen that allows either a “Full Install” or any of the individual components to be installed:

Figure 8 PFR Installer



Select any of the options and press the **Install** button, and then follow the installation instructions to complete the installation. Choose “Full Install” to perform a “typical” installation.

Installation Type Options

The PFR components installed depend on the type of installation you choose. These components are listed after you make a selection from the pull-down menu.

Full Install

The following are installed when you choose “**Full Install**”:

- PFR Control Server
- PFR Configuration
- PFR Status GUI
- PFR Index Engine
- PFR Index Configuration
- PFR Transfer Engine
- PFR Log Viewer

PFR Server Installation

The following are installed when you choose “**PFR Server**”:

- PFR Control Server
- PFR Configuration
- PFR Log Viewer

PFR Client Installation

The following is installed when you choose “**PFR Client**”:

- PFR Status GUI

PFR Indexer Installation

The following are installed when you choose “**PFR Indexer**”:

- PFR Index Engine
- PFR Index Configuration

PFR Transfer Engine Installation

The following is installed when you choose “**PFR Transfer Engine**”:

- PFR Transfer Engine

Note: During the installation process you will be asked for a PFR User Name and password. Enter the User Account information that was created during the prerequisites stage.

The user name must be preceded by the local hostname. For example, if installing on the PFRServer, the user name would be PFRServer\Administrator. Network or domain users are not supported.

Licensing your Product

Overview

The Quantum PFR suite of applications is licensed. The license enables the features of the applications. When issued, the license is tied to the MAC Address of the machine where the software was installed. Licensing applies only when you choose a Full Install or PFR Server installation.

Activating the Initial Temporary License

When the Quantum PFR Configuration Tool is first run, it automatically generates a temporary evaluation license that is valid for 30 days. During this time, Quantum PFR is fully functional, but once this 30 day evaluation period has elapsed, while the indexing process will continue, it is not possible to make any partial file retrievals.

When the initial temporary license is generated, a token file is placed onto the desktop of the Windows Server running the PFR configuration tool. The name of this token file is prefixed with the host name of the Windows Server. For example, when running the PFR configuration tool for the first time on a Windows Server called PFR-SERVER1, a token file with the following name is generated on the Windows Desktop:

`"PFR-SERVER1 PFR Product Token.txt"`

Keep this file—it is required to obtain a permanent license for Quantum PFR.

Obtaining a Permanent License

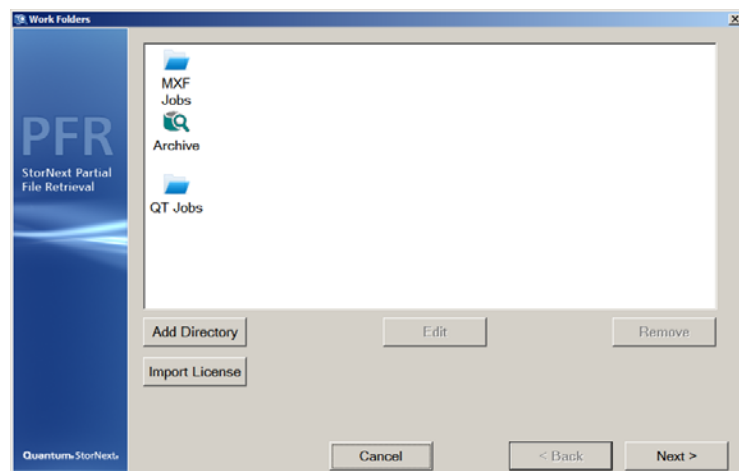
The permanent license for Quantum PFR is generated using the token file mentioned above. Your Quantum representative will send the token file to Quantum during the installation process and will receive the permanent license file back. Your Quantum representative will install this license for you.

If for any reason the permanent license is not installed as part of the Quantum PFR Installation, you will need to contact Quantum Support. The support representative will instruct you where to send your token file so you can obtain the permanent license file. The permanent license file ends with “.LIC”.

(Reminder: The token file is placed onto the desktop of the Windows Server running the PFR configuration tool. You can identify the token file because its prefix is the host name of the Windows Server. For example, if the host name is PFR-SERVER1, the token file is called “PFR-SERVER1 PFR Product Token.txt”.)

Installing the Permanent License

- 1 Copy the license file from Quantum to the desktop of the Windows Server machine that runs the Quantum PFR Configuration Tool—where the token file was initially generated.
- 2 Launch the Quantum PFR Configuration Tool.

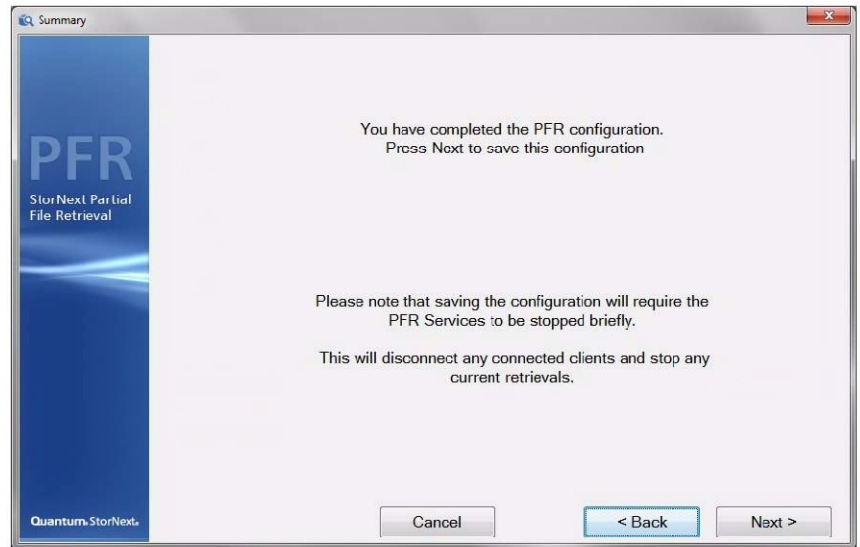


- 3 Click **Next** to navigate to the Work Folders page.

- 4 Click **Import License**. Locate the license file on the desktop and open it.

The license file is imported and validated against the MAC address of the machine to ensure that it matches the original token file.

- 5 Continue clicking Next until you reach the Summary Screen where you can save the configuration.



Saving the configuration momentarily stops the Quantum PFR Services; when they are automatically restarted, they will utilize the license imported above.

Extending a Temporary License

Under certain circumstances you can extend a temporary license for an additional 30 days, for example when needing to evaluate Quantum PFR for a greater period of time. Contact your Quantum representative to request this if required.

The process for obtaining an extension to the temporary license is identical to that for obtaining a permanent license— the token file must be sent back to Quantum and a license file will be returned. The extension license is imported into the Quantum PFR Configuration Tool in the same way as a permanent license.

Migrating a License to a New Machine

The Quantum PFR license is tied to the hardware upon which the token file was initially generated. If you need to move a Quantum PFR installation to new hardware, *do not* attempt to import the original license into the new installation as this will not work. Follow the steps below to migrate the license to a new machine:

- 1 Proceed with the software installation on the new hardware as normal.

A token file is generated.

- 2 Forward token file to your Quantum representative, letting the representative know that this is a hardware migration.

A new license file is returned to you.

- 3 Import using the [Installing the Permanent License](#) procedure.



Chapter 4 Configuration

Most of the Quantum Partial File Retrieval configuration is handled via a single configuration tool (the Primary Configuration). In addition, since the PFR Indexer can be installed on a separate machine, this element is configured separately (although it can be called from the main configuration application if both are installed on the same machine).

Virtual Machine From Earlier Version of Quantum PFR

If you have used previous versions of Quantum PFR you will be familiar with the PFR Linux Component which ran inside a Virtual Machine. Communication with StorNext is now performed using the StorNext Web Services API, and therefore the Virtual Machine is no longer required.

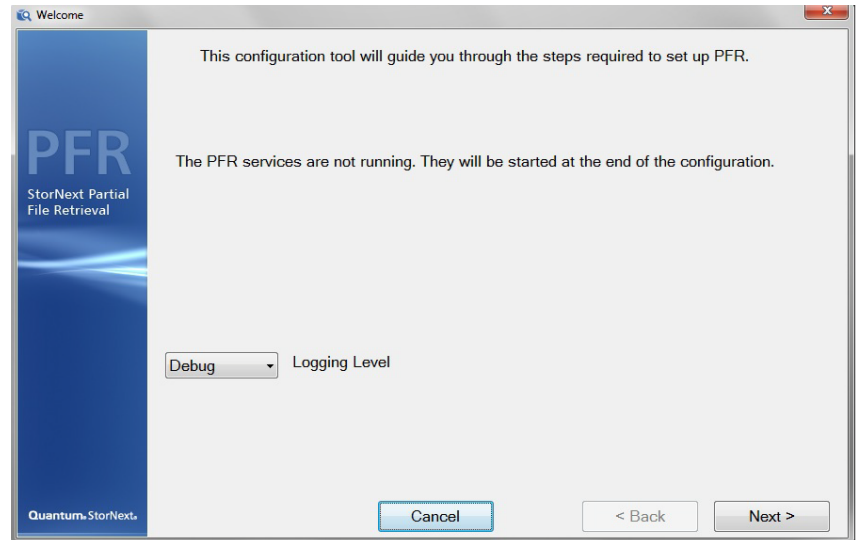
Primary Configuration

The main wizard-based tool guides you through the process of setting up the application. Configuration should only take a few minutes.

Welcome

Use the Welcome screen to select the server's logging level.

Figure 9 Welcome Screen



Press **Next** to continue.

For more information about logging, refer to [PFR Log Viewer](#) on page 57.

Archive

The Archive pane allows you to enter the machine on which the storage resides (the archive host) and the SNFS share on this storage that will be used for access to this storage. These can either be typed in manually or, better, looked up using the "Browse..." buttons.

Note: The StorNext Share should be entered as the top level share as shown in [Figure 10](#).

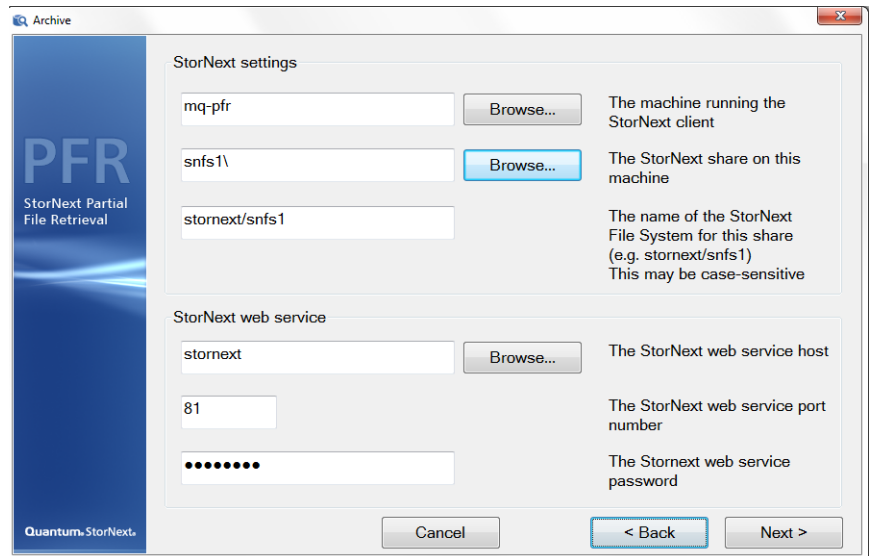
Enter the path where the file system has been mounted on the StorNext system (e.g., stornext/snfs1).

The StorNext Web Service parameters must be entered as these are used by Quantum PFR to communicate with StorNext. Your Quantum

installation engineer will be able to tell you where the StorNext Web Service API was installed.

The default Web Services port is 81, so use that port number unless it has been changed. Enter the Web Services password you created in [Defining Your Web Services Password](#) on page 19.

Figure 10 Archive Screen



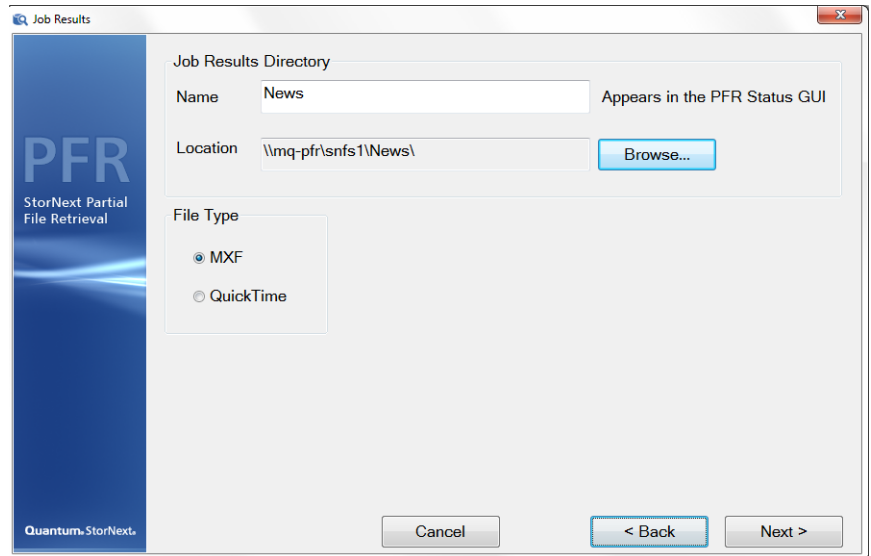
Press **Next** to continue. The StorNext Web Service connection will be tested at this point and you will not be able to proceed unless valid information has been entered.

Note: If you are upgrading from PFR 1.0.x, PFR 1.1 may not function correctly if you do not update the information in the “StorNext Web Service” section of the Archive screen. For more information, refer to [Entering Web Services Information](#) on page 17.

Job Results Folder

It is possible to create up to twenty Job Results Folders for the application, each for different roles within your organization (for example, News, Sports and Current Affairs). This page allows you to set up the first of these folders.

Figure 11 Job Results Screen



The **Name** is simply a friendly Name that is displayed in the Status GUI, while the **Location** is the UNC path to the parent location of the folder set (see [Folder Structure](#) on page 5).

Finally, select the desired File Type for this Job Result: either MXF (the default,) or QuickTime. Note that this setting will generate partial files of the specified type, irrespective of the source file format.

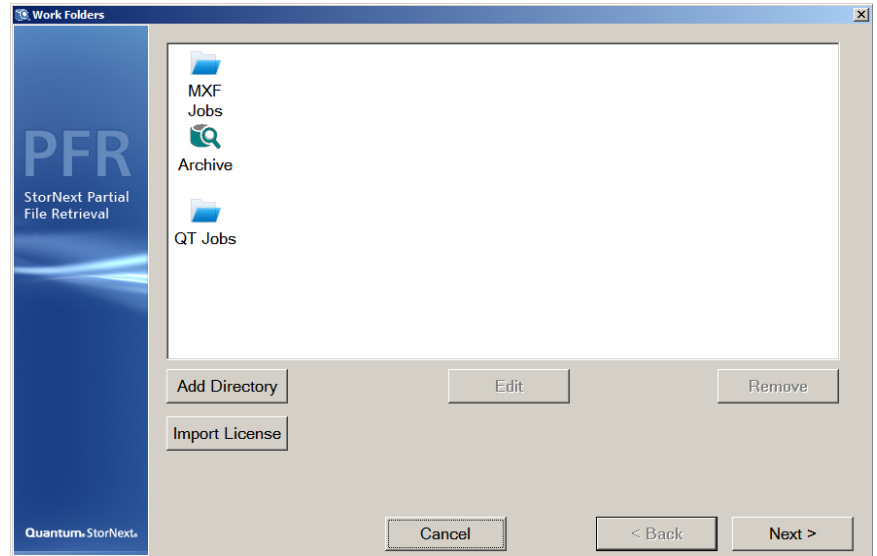
Press **Next** to continue.

Work Folders

At this stage most of the configuration has been completed. The Work Folders page provides a summary of the Archive and Job Results folders that have been set up so far. From this point you can use the **Add Directory** button to add an additional Job Results folder (up to a maximum of 20), or **Edit** or **Remove** to change or delete previously configured folders.

It is also possible to manage the Quantum PFR Licenses from here by clicking **Import License**.

Figure 12 Work Folders Screen



System Configuration

There are two aspects that are configured on this page:

- Transfer Engines - providing an ability to increase performance by running retrieval transfer engines on separate hardware;
- PFR Control Servers - providing a redundant system by replicating the PFR Control Server;

Please see [Chapter 5, PFR Active/Active Clustering](#), which addresses these topics at a conceptual level.

Transfer Engines

In a standard installation where everything is installed onto a single machine, there will be only a single Transfer Engine, located on the current machine. This will be configured automatically and no further configuration is necessary.

However, in more extensive installations where Transfer Engines have been installed on separate Windows server machines, this is the place to specify the location of the separate machines.

This window also allows you to specify the number of simultaneous transfers that should be attempted at once. The value entered here should not exceed the number of tape drives available to perform restores.

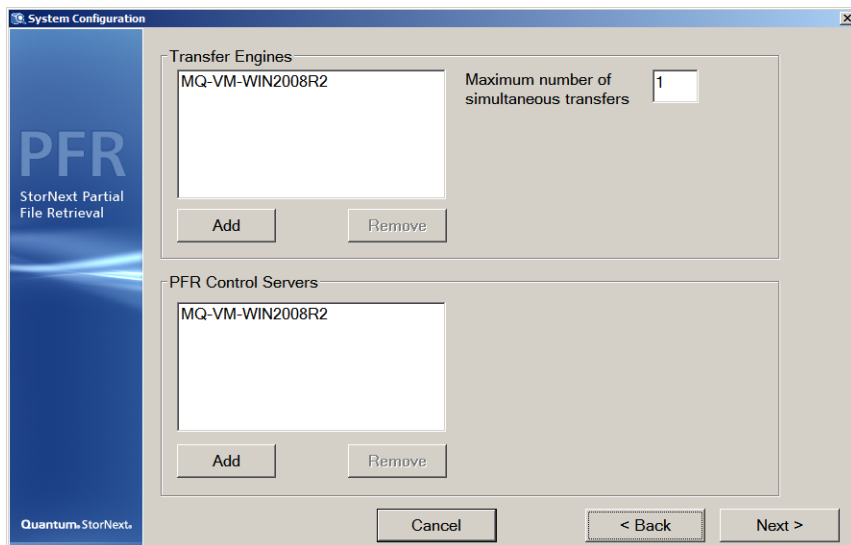
PFR Control Servers

In a standard installation where everything is installed onto a single machine, there will be only a single PFR Control Server. This will be configured automatically and no further configuration is necessary.

In order to enable PFR Active/Active Clusters, this is the place to specify the location the name of the other Windows Server machine on which the PFR Control Server software has also been installed. **PFR Control Servers should *never* be configured as localhost or 127.0.0.1.**

Please see [Making Changes in a PFR Active/Active Cluster Configuration](#) on page 42 for additional information about creating or changing the configuration in a redundant PFR Control Server environment.

Figure 13 Transfer Engines Screen



This window also allows you to specify the number of simultaneous transfers that should be attempted at once. The value entered here should not exceed the number of tape drives available to perform restores.

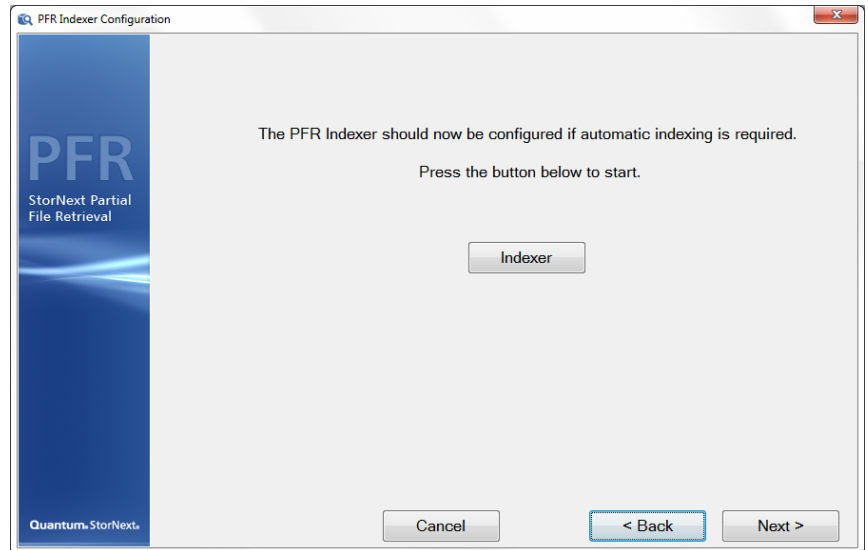
When you are ready to set up both Transfer Engine and Control Servers, press **Next** to continue.

Indexer

Configuring the Indexer is accomplished using a separate application since the indexer might be physically present on a different Windows server machine. However, in case it is present on this machine it can be launched here for convenience.

Note: The Indexer Configuration relates specifically to the ability of Quantum PFR to scan for new media files and generate index files automatically. It is not required if the PFR Indexer Web Service API will be used to create index files. Please see the next section (Indexer Configuration and Monitoring) for more details.

Figure 14 PFR Indexer Screen



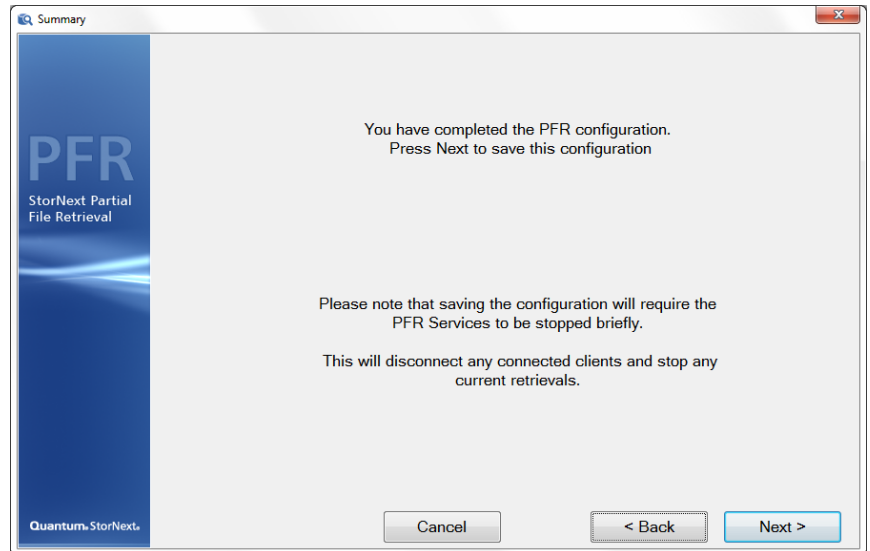
Press the **Indexer** button to launch the indexer configuration. Once complete (or if skipped), press **Next** to continue.

Summary

This page provides a final check before saving and committing the current configuration.

Caution: Saving the configuration restarts all Partial File Restore services and will therefore terminate any in-progress retrievals.

Figure 15 Summary Screen

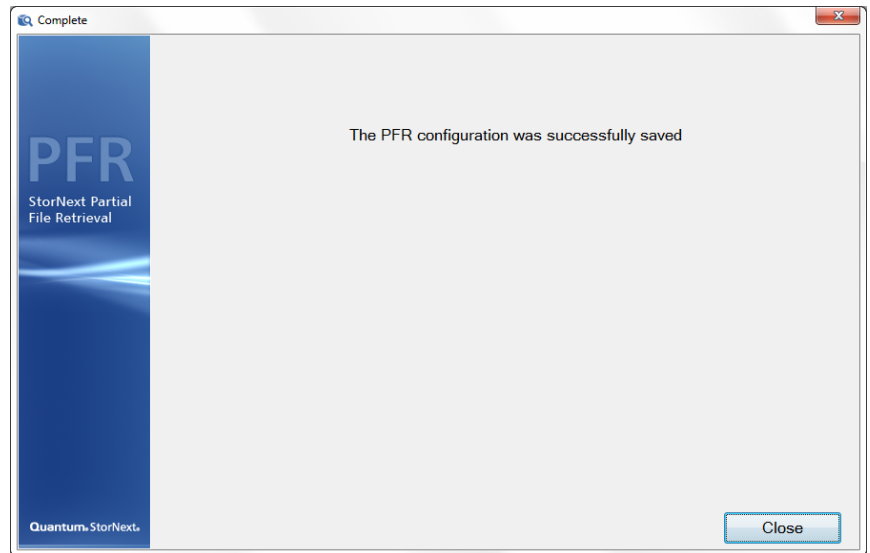


Press **Next** to continue and save the current configuration.

Complete

The final page indicates whether the configuration has been saved (or else provides details of any problems).

Figure 16 Complete Screen



Indexer Configuration and Monitoring

The Indexer is implemented as a Windows Service. The configuration and monitoring of the service are combined into this application.

Note: The Indexer service continues to run in the background, even when the PFR Index Configuration application is not running.

This application can be launched either via the main configuration tool (if they are on the same Windows server machine,) or via the Start button.

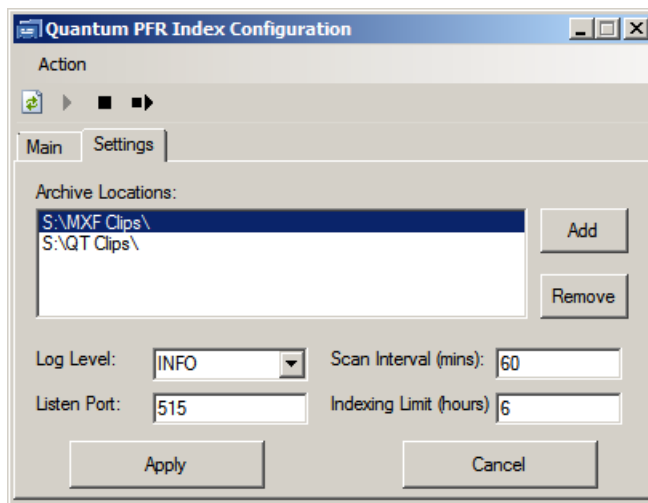
Note: This service **must** be run as a Local service.

Configuration

This application contains two tabs: one for configuration and one for monitoring. The second tab is for the configuration (settings).

Note: If the Web Services API is to be used for indexing, no further configuration of folders to scan need be entered here.

Figure 17 PFR Index
Configuration



On this screen you can do the following:

- Set one or more **Archive Locations** for which you would like to create indexes. To make changes, you can use the **Add** and **Remove** buttons on the right, or right-click with the mouse on the **Archive Locations** list. Ideally, the locations entered here will only contain media files.
- Set the **Log** (logging) **level**. (For more information about logging, refer to [PFR Log Viewer](#) on page 57.)
- Set the **Scan Interval** (in minutes). Sixty minutes is the recommended interval, but the figure is a balance between speed of indexing and the load the indexer puts onto the storage. There is no benefit in creating indexes as soon as possible as long as they are created before the files are truncated by StorNext Storage Manager.

- The **Listen Port** is the TCP/IP port that this application and the Indexer Service use to communicate on. The default TCP/IP port is 515, and this should not be changed.
- Specify an **Indexing Limit** (hours). Because there can be two or more indexers running in parallel, a locking mechanism prevents two indexers from attempting to index the same file. As soon as an indexer starts to process a clip, it “locks” the video file so that any other indexer which attempts to index the same file will see the lock, and skip the file. In the event that a machine crashes while PFR is indexing a file, a lock is applied to the file which prevents it from getting indexed by the other indexer.

The **Indexing Limit** option allows you to specify the age in hours that a lock file will be ignored so it will eventually be processed by other indexers. This is labelled “Indexing Limit (hours)” in the Quantum PFR Index Configuration application. The default value for Indexing Limit is 6 hours. As a general guideline, the limit should be less than your truncation time setting with a recommended maximum of no more than 18.

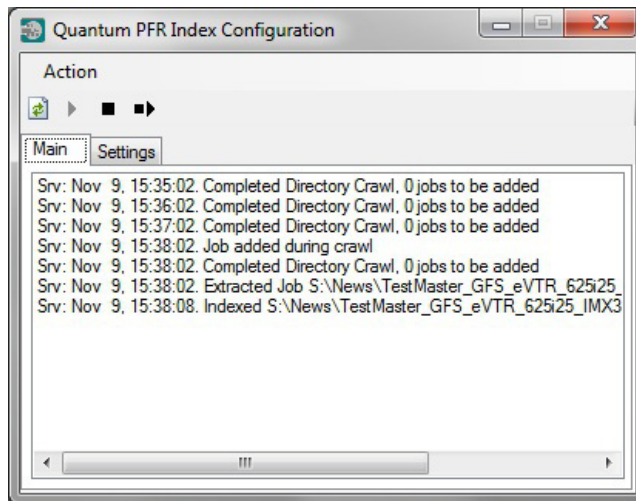
After any changes have been made, click **Apply** to save and commit your changes.

Monitoring

The Indexer tool also allows monitoring of the indexing process. As the Indexer Service scans the configured Archive Locations, log messages are written to the Log Window. At any given time you could view items indexing or items being skipped.

In a multiple-indexer environment it is normal to see “skipped file” messages, because some files are locked by one server for indexing. Then, when the other server comes across the locked files it skips indexing those files because the other server has previously claimed them.

Figure 18 Indexer Log Window



It is also possible to stop / start the Indexer Service using the toolbar buttons near the top of this screen.

Making Changes in a PFR Active/Active Cluster Configuration

When Quantum PFR is configured for PFR Active/Active Clusters, one of the two Quantum PFR Control servers will be acting as the Primary Server and one as the Secondary Server, with their databases being synchronized by the Synchronization Services. (For information about setting up PFR Active/Active Clusters, refer to [PFR Active/Active Clustering](#) on page 45.)

In order to make a change to the Primary Configuration, it is necessary to perform the following steps to ensure that the configuration is replicated correctly.

For clarity in the following description, consider a PFR Control Server configuration in which two PFR Control Servers are configured: **PFRControl1** and **PFRControl2**.

- **Manually stop the synchronization service on *PFRControl1*.** Either machine could be chosen - the effect of stopping the synchronization service on this PFR Control Server is that the other PFR Control Server will now take over as the Primary Server. For details of stopping the Synchronization Service, see the section below;
- **On *PFRControl2*, run the Primary Configuration Tool** and make any changes required to the configuration. At the end of the configuration the tool will momentarily stop the services, save the database and will then restart the services again. Because the Synchronization Service on PFRControl1 is stopped, PFRControl2 will remain the Primary Server.
- **Back on *PFRControl1*, restart the Synchronization Service.** This will link back up with PFRControl2 and accept a copy of the database which contains the configuration changes, will resume as the Secondary Server.

Starting / Stopping the Synchronization Service

The Synchronization Service is a standard Windows Service and is started and stopped through the standard Windows services list. There are many ways to bring up the list of services, one of which is to “run” the command “services.msc”:

- Press the **Windows Key + R** to bring up the Windows Run window;
- Type “services.msc” (without the quotation marks) and press the **OK** button.

The Synchronization Service is called:

MedwaySync

Once selected, this can be started and stopped using the buttons at the top of the Services List.

Note: This service **must** be run as a Local service.

Chapter 4: Configuration
Making Changes in a PFR Active/Active Cluster Configuration



Chapter 5

PFR Active/Active Clustering

Overview

Instead of installing all components on one Windows server it is possible to split the various parts of the software to obtain a level of system redundancy.

PFR Server: Includes “PFR Control Server and Configuration GUI”

- Can be installed as a single instance on a single Windows Server
- Can be installed as a redundant Control option, where a Primary and Backup Control Server will be running on two Windows server machines.

PFR Client (aka Status GUI): Allows the PFR queues to be remotely managed

- Can be installed as a single instance on a single Windows machine.
- Can be installed on multiple Windows machines.

PFR Indexer: Includes “Index Engine and Configuration / Status GUI” and the “PFR Indexer Web Service API”

- Can be installed as a single instance on the PFR Server machine
- Can be installed on its own, on a separate machine when additional performance is required.

- Can be installed on two different windows machines to enable redundant indexing in event of single machine failure.

PFR Transfer Engine:

- Can be installed as a single instance on the PFR Server machine
- Can be installed on its own, on a separate Windows machine when additional performance is required.
- Can be installed on multiple Windows machines.
 - Instances of Transfer Engine installations should not exceed the number of expected available tape drives.


PFR Server Redundancy

The PFR Server contains a small SQL-based database. This database holds the current configuration, it also contains the basic metadata for all Request EDLs it has browsed as well as the 'transfer list'.

Essentially any transfer from any source to any destination however initiated is entered into this global transfer list, which details amongst other things the name of the clip to be transferred, the source, the destination and the timecode for the media.

Therefore at a very base level the redundancy with respect to control functionality is to keep two copies of the database, one on a Primary machine, and one on the Secondary machine.

Figure 19 PFR Server Redundancy



To enable this there is synchronization service set up to run between the two PFR Servers.

Each PFR Server is aware of the other in the server set, and between them they will agree on which server acts as a Primary server and which is the Backup.

Start-up

When the physical machine is started, it starts the Synchronisation Service, this will be aware of the other Sync server in the set, and will attempt communication with it.

- 1 If it fails to communicate it will become Primary and take over that role.
- 2 If it connects, it will elect a Primary based on several internal factors.

Normal Operation

A PFR Client will only be able to connect to the Primary PFR Server.

The primary and backup have a heartbeat every 5 seconds or so to determine connectivity between them.

When the system performs a transfer, delete, or other operation the local database is updated and this new operation is mirrored by the synchronization service to the Secondary or backup PFR Server.

Shutdown

When the system needs to be shutdown, the Primary PFR server will inform the other server that a shutdown has been requested, and this server will now become the primary. To ensure that the whole system is stopped, the system needs to be stopped on both Servers.

Failure of a Primary Server

If the Secondary PFR server detects the loss of the heartbeat then it will become the Primary, and any clients connected to the initial primary if connected at the loss of the server will indicate that the Primary server has stopped and that the Client application will need to be restarted.

Any transfers physically in progress at the time of a Control Server failover will be stopped. In most cases these files must be manually resubmitted. (For information about recovery steps you should take after a failover, see [Resubmitting Requests After Control Server Failover](#) on page 51.) When client applications are restarted after a Primary Control Server failover, they will connect to the new Primary.

Note: When a Control Server failover occurs, if the Status GUI is up and running it must shut down and restarted so it can connect to the new primary system. When it does this, any jobs that had failed will be cleared so you'll see only the jobs that completed.

Failure of a Backup Server

If the Primary PFR Server detects the loss of the backup machine it will attempt, on a regular basis to re-establish communication with it.

Recovery of PFR Server

If a new PFR Server is started, the primary PFR Server is responsible for updating it to the current configuration. This will be done via a transfer of the database file wholesale.

Hardware Configurations

There are three possible configurations:

- [Combined Control Server and Engine](#)
- [Combined Control Server and Engine with an Additional Engine](#)
- [Redundant Control Server with at Least Two Engines](#)


Note: You can install the indexer (“PFR Indexer”) and/or Status GUI (“PFR Client”) separately on any of the machines in the following configurations, and/or on separate machines.

Combined Control Server and Engine

In this configuration the control server contains the Synchronisation Service, and the Engine Server Service.

One physical machine required. When configuring the (only) machine, you would choose “Transfer Engine” and then “PFR Server,” or “Full Install” if you want the indexer and client on the same machine.

Figure 20 Combined Control Server & Engine




Combined Control Server and Engine with an Additional Engine

In this configuration 2+ physical machines are required. When configuring the left machine, choose “Transfer Engine” and then “PFR Server,” or “Full Install” if you also wanted the indexer and client on the same machine.

When you configure the right machine, choose "Transfer Engine".

Note: The number of engines should not exceed the number of expected available tape drives.

Figure 21 Combined Control Server & Engine with Additional Engine




Redundant Control Server with at Least Two Engines

In this configuration at least 4 physical machines are required.

When configuring the top machines, choose "PFR Server."

When configuring the bottom machines, choose "Transfer Engine."

Figure 22 Redundant Control Server with 2+ Engines



Location of PFR Indexers

In any of the preceding configurations, greater performance and system redundancy can be achieved by running instances of the PFR Indexer on multiple Servers, often on Primary / Backup Control Servers. For PFR Active/Active Clustering, the configuration of each PFR Indexer must include all indexed directories.

Resubmitting Requests After Control Server Failover

This section describes how to recover files which either fail or don't process correctly during Control Server failover. These requests fail or do not finish processing due to a loss of connection to a Control server that may be managing the transfer.

Before starting this process the two Control Servers should both be online and have their MedwaySync service in the Started state. If this is not possible, then a reduced system with the failed Control Server completely offline (either shutdown or all PFR related services stopped) will also be acceptable. The PFR service names are

- Marquis Log Service
- Medway Engine Service
- MedwaySync
- Quantum PFR Indexer Service

See [Starting / Stopping the Synchronization Service](#) on page 43 for directions on stopping these services.

To resubmit requests that have failed or have not processed correctly, perform the following steps:

- 1 Open the PFR Status GUI. Confirm that no jobs are currently in the "In Progress", "Parsing" or "Pending" states. All requests should be in either the "Completed" state or in a failed state.

At the bottom of the right pane there should be several EDL requests in a failed state that have red messages next to them. Here are some messages to look for:

- "File locked by another process"
- "OK" (in red)

- “Could not connect to Transfer Engine to perform Transfer”
 - Any other unexpected failure-related message
- 2 Open Windows Explorer and browse to the “PFR-REQUEST” directory in the Job Results folder you're working in. This directory's location was defined during the initial PFR configuration. In this directory you will find two possible file types:
 - I. XML Files
 - a Create a new directory in the Job Results directory location. You can name it anything you like.
 - b Any EDL request file in the PFR-REQUEST directory that ends in .XML should be copied from this directory and saved into the directory created in the previous step (I.a) for use later in this process.
 - II. LCK Files
 - a Any EDL request file that ends in .LCK should be opened with a text editor such as Notepad or Wordpad. While viewing each file, note the contents of the <Title> tag.
 - b In Windows Explorer, navigate to the main Job Results directory. If there are PFR output files (.mov or .mxf) with the names that were recorded in the previous step (II.a) and the file size is increasing after doing a Windows Explorer refresh, then the .LCK file is valid and should be in the PFR-REQUEST directory as the job is still in progress.
 - c If there are no PFR output files with those names or there are PFR output files that have an unchanging file size after a Windows Explorer refresh, then the .LCK file is invalid and should be deleted.
 - 3 In Windows Explorer, navigate to the “fail” directory in the Job Results directory you've been working with in Step 2. Here you will find all your failed EDL request files. Select any EDL request files with a date created (in Windows Explorer) of the time the retrievals were processing. This can be determined by looking at the “Completed” time on successful requests in the PFR Status GUI. Copy the EDL request files to the directory created in step II.a.
 - 4 Returning to the PFR Status GUI, clear the right pane of entries by selecting all request entries (success and failure) and clicking on the trash can in the lower right corner of the Status GUI. Only select the “Delete reference movie” option and click “OK”.

- 5 Copy all EDL request files saved in steps 1.b and 3 into the PFR-REQUEST directory in the Jobs Results directory you're working in.
- 6 Monitor the Status GUI for request results.



Chapter 6

Status and GUI Application

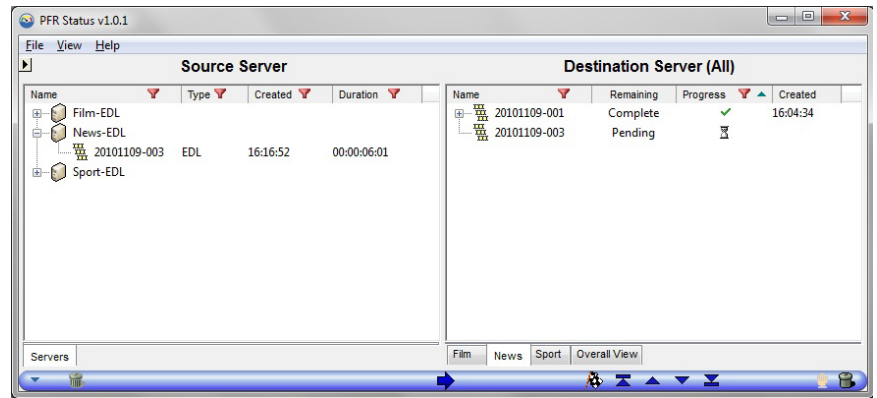
By default, the PFR Services should handle all retrieval requests automatically. However, should you wish to monitor the process, manipulate the pending requests or look at the log files, this can all be done using the Status GUI application.

Anatomy of Status GUI

The left side of the Status GUI screen lists the PFR source requests, which are the XML request files that have not yet been processed, categorized according to the folder sets defined in the configuration.

On the right are the destinations: one tab (along the bottom) for each configured folder set. The entries on the right represent pending, in-progress and complete (or failed) retrievals.

Figure 23 Status GUI Screen



Manipulating the Retrievals

The following actions can be performed on retrievals using the Status GUI. In the following descriptions, any action will reflect the selected transfer.


- **Prioritization.** Using the set of five buttons below the **Destination Servers**, a pending transfer can be given high or lower priority by moving the pending transfer to the top, up, down or bottom of the list. All high priority transfers will take place before all standard priority transfers.
- **Delete.** Using the trash icon in the bottom right, pending transfers can be removed. That also cancels active retrievals and deletes any complete transfers (it will remove the partial file).
- **View Log File.** From the **Help** menu it is possible to view the current log file. If the Status GUI is run on the same machine which runs the main PFR services, this will be the log file which also contains all of the retrieval information.

PFR Log Viewer

To be able to read the log file from the PFR Status GUI, there is an additional application that is installed with a full PFR server installation. This application is the PFR Log Viewer, which enables a user to merge logs from different machines running components of the PFR system.

The application can be launched from Start/All Programs/Quantum PFR/PFR Log Viewer.

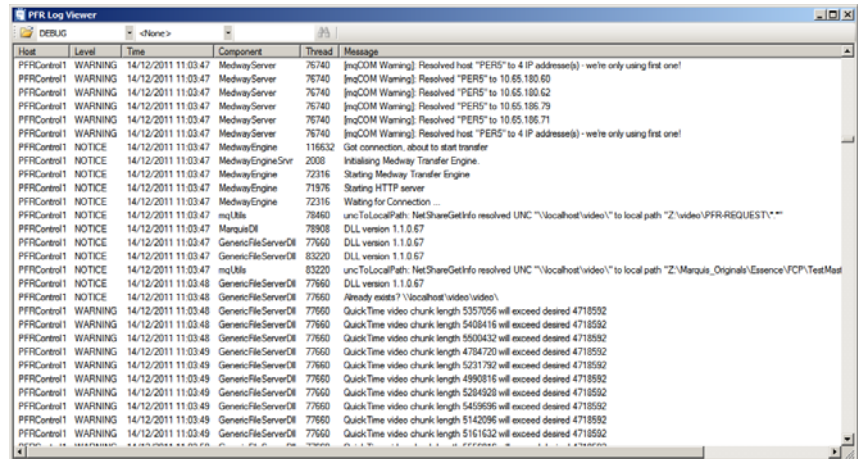
Figure 24 Blank Log Viewer Screen



After launching you will need to open the desired log file, which is achieved by using the 'Folder Open' button (top left of the GUI), and navigating to the log.

Note: Logs are located on all machines where PFR is running in the C or D drive (depending on the installation drive selected) /Program Files (x86 for 64-bit machines)/Quantum/Logs. In here the log files are named using a DD-MM-YYYY format, as a new log file is created every 24 hours.

Figure 25 Populated Log Viewer Screen



The level of the logging can be selected by choosing a range selection between Notice and Debug.

Following is a summary of available logging levels:

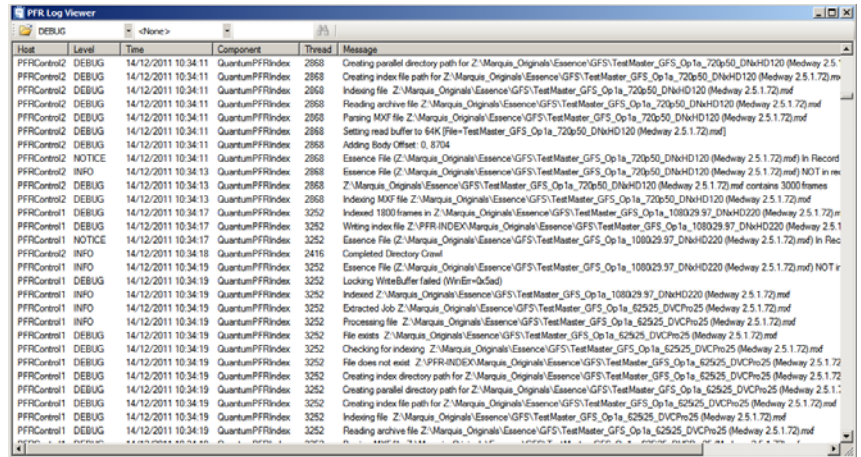
- **EMERGENCY/ALERT/CRITICAL** typically means that part of the system or the whole system will stop behaving normally, or may quit altogether.
- **ERROR** usually means that a job will fail.
- **WARNING** gives clues as to why the job failed, but may not result in a failure.
- **NOTICE/INFO** gives verbose details about which operation is running. This is the default logging level.
- **DEBUG** lists in great detail everything the system is doing. Large logs may be created, but this option is often very useful to determine the exact point at which an operation error occurred, and why.

The log above is set to DEBUG display and next to this is a selection for thread number, the display shows this a none. Thread numbers are important on a busy system as you will be able to track a single restore more easily. (For example, if two restores are happening simultaneously they will each have a unique thread number.)

To merge multiple log files from different physical machines running PFR you only need to open an additional log file using the open folder

button. Please be aware that you will need to select a same dated log file.

Figure 26 Merged Log Viewer Screen



Once multiple log files have been opened, the Host column will show the machine name of entries that came from remote machines, or the drive letter for entries that came from the local machine.



Chapter 7

Request EDL Format

The following example Request XML file should provide sufficient details to allow for creating requests.

Quantum does not recommend “hand creating” XML files. XML files should be created by a MAM or similar application.

Sample

```
<?xml version="1.0" encoding="utf-8"?>
<MarquisEDL version="2.0">
  <FrameRate>DF</FrameRate>
  <ClipList>
    <Clip>
      <UID>ID1234567891</UID>
      <Title>MyFirstPFR</Title>
      <File>//MyServer/SNFS1/Dir1/Dir2/MySourceClipFile.mxf</File>
      <Start>10:01:30;10</Start>
      <End>10:02:15;09</End>
      <NumberVideoTracks>1</NumberVideoTracks>
      <NumberAudioTracks>4</NumberAudioTracks>
    </Clip>
  </ClipList>
</MarquisEDL>
```

Notes

FrameRate	"25"	: PAL
	"DF"	: NTSC Drop Frame
	"NDF"	: NTSC Non Drop Frame
	"50"	: 50 frames per second
	"60NDF"	: 60 frames per second
	"60DF"	: 59.94 frames per second
	"24NDF"	: 24 frames per second
	"24DF"	: 23.976 frames per second
Clip	There can only be one clip provided	
UID	A unique ID for the source clip, but otherwise not used	
Title	Name of the (partial) output file	
File	The full UNC path to the source file on the StorNext managed storage	
Start	Start timecode of partial clip to be retrieved	
End	End timecode of partial clip (inclusive - e.g. timecode of last frame of partial clip)	
NumberVideoTracks	Number of video tracks in source clip. Should always be "1"	
NumberAudioTracks	Number of audio tracks in source clip	

Special Characters

The **Title** field above is used to generate the file name of the resulting partial file. Therefore, ensure that the value put into this field is a valid file name.

Note: Quantum PFR runs as a set of Windows Services. Therefore, the file name must be valid both as a Windows file name and as a valid file name on the managed disk.

The **Title** field must not include any the following characters:

Colon	:
Backslash	\
Forward Slash	/
Asterisk	*
Question Mark	?
Double Quotation Mark	"
Less than	<
Greater than	>
Pipe (or Bar)	



Chapter 8

Performance and Scalability

The performance of Quantum PFR will be affected by a number of factors, and the interactions are complex. This chapter provides some guidelines, but actual performance depends largely upon the environment in which it is run.

Performance

Indexer

The Quantum PFR Indexer usually starts to index a file within two or three hours of it being written, although the StorNext policy should be set to truncate files no sooner than 24 hours.

The performance of the indexer is influenced by a number of factors, including:

- The file wrapper type - MXF files require a greater effort to index than QuickTime files
- The size and duration of the files. Of course, longer files will take longer to index, but a longer file can be indexed more efficiently than a number of smaller files of the same total size / duration.

Based upon the recommended hardware listed in [Suggested Platform](#) on page 69, tests have shown that the Quantum PFR Indexer can index a

24GB MXF file in approximately 6 minutes. This would equate to 5.4 TB or 120 hours of HD material at 100 Mb/s in a 24 hour period.

Note on Existing Archives

There is no requirement to index files before they can be retrieved by Quantum PFR. Where no index file yet exists, Quantum PFR retrieves the whole file from tape and then extracts the required partial file from the whole file.

This whole file would then be indexed before it was truncated again, meaning that the index would be available for subsequent requests. While this is less efficient than a strictly partial retrieval from tape, it does allow Quantum PFR to begin to operate on an existing archive straight away.

Partial File Retrieval via Index from Tape

Like the Quantum PFR Indexer, there are a number of factors that affect the performance of the retrieval from tape:

- There is a fixed overhead for the tape handling, working on worst case assumption that each retrieval requires unloading the previous tape and loading the new tape. A slightly conservative estimate value of three minutes is assumed for this process.
- There is an additional overhead of approximately 30 seconds for Quantum PFR to submit the request to Storage Manager, detect the completion of the partial file, generate headers for the partial file and tidy up at the end of the transfer.
- Retrieval from Tape to Spinning Disk (which is very fast).
- Subsequent creation of the partial file.

Taking into consideration the overheads, the average sustained data rate for the re-wrapping the partial file (averaged over 24 hours) would be no less than 100 Mb/s.

Very Short Partial Retrievals

With very small retrievals, the tape handling overhead tends to be the largest part of the time taken to perform a partial file retrieval. In this case, the total capacity of the system tends towards being limited by the total number of retrievals requested. Assuming approximately four minutes of overhead per transfer, when the transfers are very short

there will be a limit of approximately 15 transfers per hour or 360 transfers per day, per tape drive in the system available for restores.

Longer Partial Retrievals

With longer retrievals where the tape handling overhead is smaller than the time to create the partial file, it should be possible to drive the system in such a way that the tape system is kept busy getting ready for the next restore, while the previous partial file creation is being performed. Therefore, with longer partial retrievals, the limiting factor becomes the network bandwidth averaged over 24 hours (100 Mb/s).

Based upon a 100 Mb/s averaged network bandwidth, the point at which the tape handling overhead becomes smaller than the partial file creation would be approximately 3GB for the partial file.

This would equate to approximately 1TB of material being retrieved within 24 hours, which would equate to approximately 20 hours of HD material at 100 Mb/s.

100Mb/s is a very conservative estimate based on averaging the total system performance over a 24 hour period, assuming a mix of long and short retrievals. The peak re-wrapping performance would be expected to be several times faster than this average, showing a relative improvement in total performance for longer retrievals.

File and Folder Structure Limitations

Quantum PFR Indexer

The Quantum PFR Indexer maintains a memory image of the indexed folder and file structure during the indexing process. This will scale to millions of files and folders.

Note that the implementation of the Quantum PFR Indexer is geared around there not being a huge number of files per folder (less than 5,000). A very large number of files in each folder is more processor intensive to index than the same number of files, spread between a larger number of folders and will have a small detrimental impact upon indexing performance.

Quantum PFR Server

The Quantum PFR Server does not hold information about the files and folders to be restored, and is therefore unaffected by the size of the archive.

Scalability

Quantum PFR performance for both the Indexer and the Retrieval Process is determined largely by the available network bandwidth.

Standard Windows Server monitoring tools (such as Resource Monitor in Windows 2008 Server R2) can be used to establish whether the storage connection bandwidth is being fully utilized a large percentage of the time.

There are several ways in which the Quantum PFR system can be expanded in order to enhance performance.

Distributing the Load

- Installing the Quantum PFR Indexer onto a separate machine separates the indexing network bandwidth from the partial file retrieval network bandwidth. This would be beneficial if there is a large amount of new material being added to the system.
- Installing an additional Quantum PFR Transfer Engine on separate hardware also spreads the load of the retrieval bandwidth. This would be beneficial when there is a large number of partial retrievals being performed.

Dedicated Tape Drives

As long as network bandwidth is not a limiting factor and retrievals are from different tapes, greater throughput of the Quantum PFR system will be obtained by adding more Tape Drives to the system, dedicated to retrievals.

Suggested Platform

- Windows 2008 Server R2 (64 bit)
- ≥ 2 CPUs with ≥ 2 GHz clock speed
- ≥ 4 GB RAM
- ≥ 8 GB RAM for archive of $> 1,000,000$ files
- ≥ 20 GB available local disk space after the pre-requisites have been loaded
- ≥ 2 Gb/s Fibre-Channel connection to SAN
- Ethernet NIC for PFR status GUI access



Chapter 9

Index Web Service API

Quantum PFR Indexing can be performed one of two ways:

- Automatically, by configuring the Quantum PFR Index Service to scan for new media files and to automatically generate the index files as they are found;
- Manually, but making calls to this Index Web Service API in order to create each index file as required.

There is no “right” or “better” way to perform indexing, the option chosen will depend upon the nature of the workflow in your organization and whether you have the development resources available to program against a Web Service API such as this.

Within a single installation of Quantum PFR, Quantum recommends that all indexing is performed by one of the two methods above and not with a mixture.

Regardless of which method of indexing is utilized, the functionality is provided as part of the Quantum PFR Index Service (a standard Windows Service application).

Note: A query to a pre-PFR 1.1 index shows only whether the file successfully indexed, and no additional information is provided. If you require extended information, delete the pre-PFR 1.1 index and re-index the file.

API Detail

The Quantum PFR Index Web Service API is a REST Interface that listens on TCP/IP port 60786. There are two commands:

Indexing a File

Call Name	indexfile
Call Type	GET
Parameter	filepath The full path (via a mapped drive on the StorNext MDC) to the media file to be indexed.
Returns	See Returned XML on page 74

Notes: This method will block while the index is created and will only return when either the index file has been created or for some reason it has not been possible to create the index file.

If a request is made to create an index for a media file which has been truncated by StorNext, this call will cause the entire media file to be restored by StorNext.

The Web Service will support multiple concurrent calls to this command.

Example

The following example shows the HTTP request to send PFRIndexer a request to index the media file S:\Test\File1.mov:

```
http://PFRIndexer:60786/  
indexfile?filepath=S:\Test\File1.mov
```


Indexing Offline Files

If the indexer is instructed by the Web service to index an offline file, it will attempt to do so and bring the file back online.

Indexing Multiple Files

If you index multiple files using the Web service, it repeatedly puts the following line in the log for as long as the lock mechanism is being used (which might be for a few minutes):

```
WARNING    : Oct 14, 16:55:55.    QuantumPFRIndex [01432] :
Locking
ReadBuffer failed (WinErr=0x5ad)
```

Requesting File Index Status

Call Name	filestatus
Call Type	GET
Parameter	filepath The full path (via a mapped drive on the StorNext MDC) to the media file whose status is to be retrieved.
Returns	See Returned XML on page 74

Notes: This method will block while retrieving the index status for a previously indexed file.

This method internally uses an XML file, generated by the indexer to retrieve the detailed status. This XML file is only generated by Quantum PFR version 1.1 and later. For files that were indexed in an earlier version of Quantum PFR, minimal information will be retrieved (just whether the file had been indexed or not).

The Web Service will support multiple concurrent calls to this API call.

Example

The following example shows the HTTP request to request from server PFRIndexer the status of media file S:\Test\File1.mov:

```
http://PFRIndexer:60786/  
filestatus?filepath=S:\Test\File1.mov
```

Returned XML

The return from either of the above calls is a packet of XML that describes the result of the call. In effect, internally, when requested to perform an index on a file, we then request its status and make that the return packet of the call.

Failed To Index

```
<?xml version="1.0" encoding="utf-8" ?>  
<IndexerReport IndexResult="Failed" IndexTime="2011/10/21  
15:30:15"/>
```

Good File Index or File Status Call

```
<?xml version="1.0" encoding="utf-8" ?>  
<IndexerReport IndexResult="Succeeded" IndexTime="2011/10/  
21 11:40:53" FileStartTC="01:00:00;00" FileDuration="1800"  
FileFrameRate="29.97"/>
```

File Status When File Not Present

```
<?xml version="1.0" encoding="utf-8" ?>  
<IndexerReport IndexResult="Error File Not Found"/>
```

File Status When File Not Indexed

```
<?xml version="1.0" encoding="utf-8" ?>  
<IndexerReport IndexResult="Not Indexed"/>
```