FUSION PACS™

Storage and Distribution Manager Implementation Guide

Version 1.50

MERGE™ Healthcare

Part number 0911-237-0001-000

April 5, 2007

Copyright notice

Copyright © 2007 Merge Healthcare
6737 W. Washington Street, Suite 2250, Milwaukee, WI 53214-5650, USA.
Phone: (414) 977-4000 Fax: (414) 977-4200

All Rights Reserved. This manual shall not wholly or in part, in any form or by any means, electronic, mechanical, including photocopying, be reproduced or transmitted without the authorized, written consent of Merge Healthcare.

This guide has been produced to assist in providing instruction for the Merge Healthcare product suite. Every effort has been made to make the information in this guide as accurate as possible. The authors of Merge Healthcare shall have neither liability nor responsibility to any person or entity with respect to any loss or damages in connection with or arising from the information contained in this guide.

Trademarks

Merge Healthcare, Fusion PACS, and Fusion eFilm are all registered trademarks of Merge Healthcare and its subsidiaries.

.NET Framework, DirectX, Internet Explorer, Microsoft, SQL Server, Terminal Client, WebControls, and Windows are trademarks or registered trademarks of Microsoft Corporation.

Application Storage Manager (ASM) and Sencor are products of Software Engineering Corporation.

DLTtape is a registered trademark of Quantum.

Linux is trademarked by Linus Torvalds.

Pentium is a registered trademark of Intel Corporation.

RealVNC is a trademark of RealVNC ltd.

Any additional software products named in this document are claimed as trademarks or registered trademarks of their respective companies.

Table of Contents

Installing Fusion PACS 1

The Fusion PACS Storage and Distribution Manager is a DICOM image archiving and distribution system. It consists of a Web server, a set of Windows NT services, and a database. These can be installed on a single machine or distributed among several machines.

This chapter provides basic instructions for installing Fusion PACS on a server.

Note: If you are upgrading from an existing Fusion PACS installation (1.4.2 or later), see Appendix C, "Upgrading From 1.4.2" on page 99.

Once you have installed the software, you still must configure Fusion PACS as described in Chapter 2, "Configuring Fusion PACS" before you can use the PACS.

Note: Fusion PACS can also be configured to operate in a clustered environment. See Chapter 2, "Configuring Fusion PACS" for more information about setting up FUSION PACS on a server cluster.

To install Fusion PACS (overall procedure):

1. Review the pre-installation procedures (see "Before you install" on page 5)
2. Install any required third-party software (see "Installing third-party software" on page 6).
3. Install the PACS software (see "Installing the Fusion PACS software" on page 8).
4. Install services and databases (see "Completing the installation" on page 12).

Before you install

Before you start the installation, you should do the following:

- If the site is using key images, make sure that the Key Images server is set up and running.
- Set up a "services" user on the machine where the Fusion PACS NT services will be running. The user must have **administrator rights** and be able to **log on as a service**.
- Gather the information needed for the installation. You can use the following table to note the information:

Required Field	Value
SQL server	
IP address	

Required Field	Value
Admin username	
Admin password	
FUSION Services user account/ password	
Workflow server *depending on configuration	
IP address	
Port number	
Report Server *depending on configuration	
IP address	
Port number	
Report URL *depending on configuration	
Key Images Server *depending on configuration	
IP address	

Step 1: Installing third-party software

Fusion PACS requires a number of third-party software products to be installed on your server before you can proceed with a Fusion PACS installation. This section describes how to determine the third-party software requirements for your installation type.

The following software is required by all installation types:

- Windows 2000 Service Pack 4 or Windows 2003 Server Service Pack 1
- Internet Explorer 6 Service Pack 1

The following software is optional for all installation types:

- Tunnel
- Terminal Client
- VNC

There are four possible software configurations for Fusion PACS. These are listed below with a summary of the third-party software required for each installation:

- **Complete**
  - SQL Server Service Pack 3a (see note below for more information)
  - .NET Framework 1.1 (if using Windows 2000 SP4)
  - ASP.NET (if using Windows 2003 Server)
  - IE WebControls

  - J#
- **Database only**
  - SQL Server Service Pack 3a
- **Services only**
  - J#
- **Web only**
  - .NET Framework 1.1 (if using Windows 2000 SP4)
  - ASP.NET (if using Windows 2003 Server)
  - IE WebControls

The following table provides a short description of each third-party software product:

Windows 2003 Server Service Pack 1	The Windows 2003 Server operating system.
Windows 2000 Service Pack 4	Windows 2000 Service Pack 4A contains a collection of fixes in the following areas: security, application compatibility, operating system reliability, and setup. Windows 2000 Service Pack 4 is a required update that includes the updates from previous Windows 2000 Service Packs.
SQL Service Pack 3a	SQL Server Service Pack 3a (SP3a) addresses specific issues discovered in SQL Server 2000 since its original release. Because SQL Server Service Packs are cumulative, SP3a includes all fixes from previously released service packs, and can be applied to an original installation or to an installation where any other service pack was previously applied.
Tunnel	Tunnel is monitoring software used by Merge Healthcare Service, used to set up a secure connection to a Fusion PACS.
.NET Framework 1.1/ASP.NET	If you are using Windows 2000, install .NET Framework 1.1; if you are using Windows 2003 Server, install ASP.NET. These are installed by the Fusion PACS installer.
Web Controls	Microsoft Internet Explorer Web Controls is a collection of ASP.NET server controls.
Internet Explorer 6 SP1	Internet Explorer 6 SP1 is a set of core technologies used in Microsoft Windows operating systems.
Terminal Client	The Terminal Services Advanced Client (TSAC) is a Win32®-based ActiveX® control (COM object) that can be used to run Terminal Services sessions within Microsoft Internet Explorer.
J#	The Microsoft J# runtime is required for any PACS install that includes services. It is installed by the Fusion PACS installer.
RealVNC	RealVNC™ is remote control software that allows you to view and interact with one computer (the server) using a simple program (the viewer) on another computer anywhere on the internet.
Internet Information Services (IIS) 6.0	Internet Information Services (IIS) 6.0 is a powerful Web server that provides a highly reliable, manageable, and scalable Web application infrastructure.

To install third-party software:

1. Insert the Fusion PACS CD into the CD-ROM drive. If Autorun is not enabled, explore the CD-ROM and run **launch.exe**. The Fusion PACS installation program begins and the Fusion PACS splash screen appears.
2. Click **Resources**. The *Resources* screen appears.
3. Select the appropriate third-party software for your intended installation type.

Note: If you are using Windows 2003 Server, the ASP.NET software cannot be installed using the *Resources* screen. Instead, you must follow the instructions below. If you are using Windows 2000, the installer will check for ASP.NET and offer to install it if it's missing.

To install ASP.NET in Windows 2003 Server:

1. In Windows 2003 Server, select **Start > Settings > Control Panel > Add/Remove Programs**. The *Add or Remove Programs* application appears.
2. Click **Add/Remove Windows Components**. The *Windows Components Wizard* appears.
3. Select **Application Server** and then click **Details**. The *Application Server* dialog box appears.
4. Select the **ASP.NET** check box.
5. Verify that ASP.NET, Enable Network COM+ Access, and Internet Information Services (IIS) are selected.
6. Click **OK**.
7. Click **Next** to update Windows 2003 Server.

Note: You may need your Windows 2003 Server CD.

8. Click **Finish**.

Step 2: Installing the Fusion PACS software

1. Insert the Fusion PACS CD into the CD-ROM drive. If Autorun is not enabled, explore the CD-ROM and run **launch.exe**. The Fusion PACS installation program begins and the Fusion PACS splash screen appears.
2. Click **Installs**.
3. Under **Fusion PACS**, click **Version 1.5**. The J# installer will launch.
4. Follow the prompts to install J#. Once the J# installation is complete, the FUSION SD installation will proceed. The *Welcome* screen appears.
5. Click **Next**. The *License Agreement* screen appears.
6. Click **Yes** to accept the license agreement. The *Information* screen appears, displaying important information about the current release of Fusion PACS SD.
7. Click **Next**. The *Customer Information* screen appears.
8. Complete the **User Name** and **Company Name** fields according to customer preference. Select **Anyone who uses this computer (all users)**, and click **Next**. The *Select Features* screen appears.

Note: By default, the options for a full installation are selected. Unless you are setting up a distributed install, or want to install the various Fusion PACS components on different physical machines, you should accept these settings.

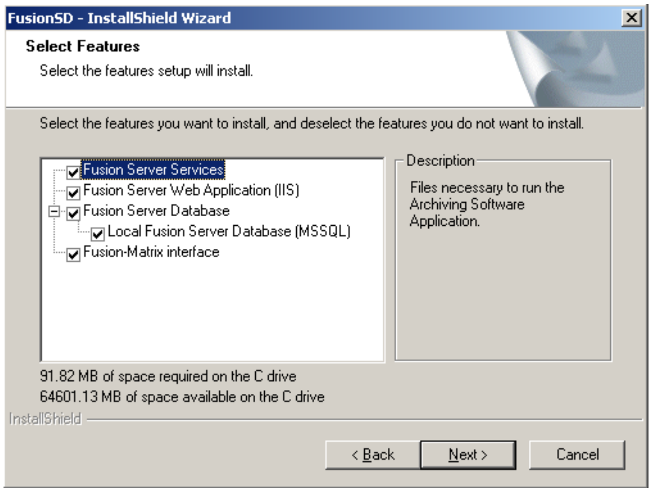

FusionSD - InstallShield Wizard

Select Features

Select the features setup will install.

Select the features you want to install, and deselect the features you do not want to install.

- ☑ Fusion Server Services
- ☑ Fusion Server Web Application (IIS)
- ☑ Fusion Server Database
  - ☑ Local Fusion Server Database (MSSQL)
- ☑ Fusion-Matrix interface

Description

Files necessary to run the Archiving Software Application.

91.82 MB of space required on the C drive

64601.13 MB of space available on the C drive

There are four feature options for the Fusion PACS installation:

FUSION Server Services	Fusion PACS NT services are to be installed.
FUSION Server Web Application (IIS)	Fusion PACS Web services are to be installed.
FUSION Server Database	Fusion PACS database component is to be installed.
Fusion Matrix Interface	Interface software to Fusion Matrix workstation. If you install this feature, you must follow the procedures described in Chapter 3, "Configuring the Matrix Interface" to configure the interface.

9. Select the features that are to be installed, and click **Next**. The *Fusion PACS - Features* screen appears.
10. Verify the services that are to be installed, and click **Next**. The *Active SQL Servers* screen displays a list of SQL Servers that are available for connection.

11. The list of servers is shown for your information only; if the server to which you intend to connect is not shown, make sure the server is running. Click **Next**. The *MSSQL Server Information* screen appears.

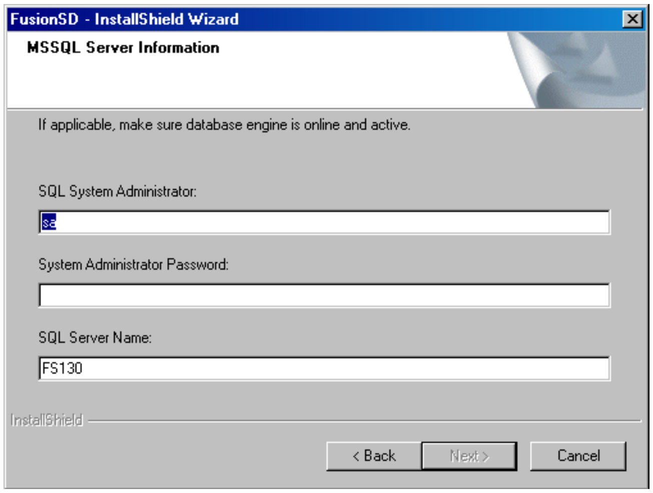

FusionSD - InstallShield Wizard

MSSQL Server Information

If applicable, make sure database engine is online and active.

SQL System Administrator: sa

System Administrator Password:

SQL Server Name: FS130

Important: You must enter the correct information in the *MSSQL Server Information* screen to proceed with installation. If the SQL Server is local, it must be accessible and online for the installation to continue. Fusion PACS verifies the entered values by querying the SQL Server. If you entered incorrect values, the system will notify you and return to this screen with default values.

12. Enter the information for the server that will contain the database, and ensure that the database engine is online and active. Click **Next** to continue. The *Database Location Question* screen appears.
13. Select one of the locations to install the database components, and click **Next** to continue.
    - If you selected **Local Directory Assumed**, continue to step 14.
    - If you selected **Custom Directory Chosen**, go to step 15.
14. The *MSSQL assumed local path definition* screen appears. Verify that the Destination Folder is correct, then click **Next** and continue to step 16.

Important: The directory for the Fusion PACS database must have 900MB of unused disk space before the paths are defined through this window. System failure in locating the space requirements will result in the system returning to the *MSSQL assumed local path definition* screen.

15. The *MSSQL custom path definition* screen appears. This screen allows you to determine where the various database components will be installed.

    Although the illustration below shows all of the database files being installed on the same drive, Fusion PACS database performance can be significantly improved when the three sets of database components are located on separate drives. You should also consider the location of the operating system page file (also known as a swap file) when determining which drive to use for the database files. We do NOT recommend placing all three database files and the operating system swap file on the same spinning volume (physical drive).

    Also, be aware of which drives are simply virtual drives (all on the same spinning disk) and which drives actually reside on separate spindles. At a minimum, the database files should at least reside on a spinning volume separate from the operating system swap file.

One typical approach is to place the operating system, its swap file and the SQL log files on mirrored drives (RAID 1), and then deploy the main database files and database index files on a separate RAID 5 array.

Note that there is no "one size fits all" scenario and the exact configuration must consider the redundancy, scalability, performance requirements, and budget of the site in question.

Important: The directories for the base, index, and log files for the database **must** be created with the required disk space allocations before the paths are defined through this screen. If the system fails to locate the paths, or insufficient space has been allocated to each directory, the installer will return to step 14 and use the default installation directories.

Enter the paths for the following sets of database files, and click **Next**.

- **Base** (500MB disk space requirement): the database files themselves
- **Index** (150MB disk space requirement): the database indices
- **Log** (150MB disk space requirement): the database log files

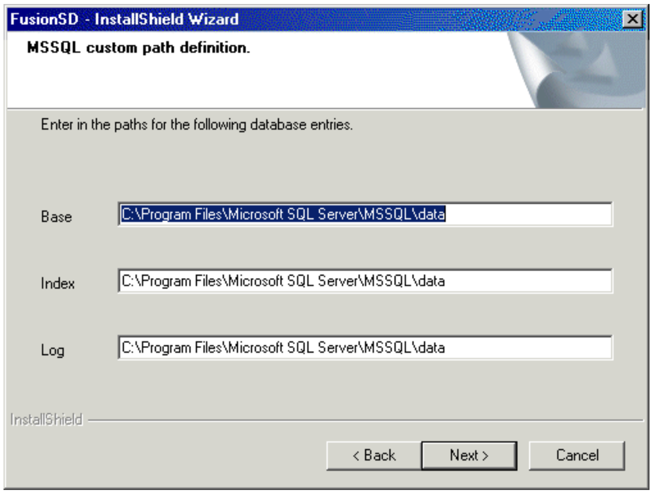

FusionSD - InstallShield Wizard

MSSQL custom path definition.

Enter in the paths for the following database entries.

Base C:\Program Files\Microsoft SQL Server\MSSQL\data

Index C:\Program Files\Microsoft SQL Server\MSSQL\data

Log C:\Program Files\Microsoft SQL Server\MSSQL\data

16. The *Choose Destination Location* screen for the Data Directory appears. This is the location where all of the customer's unarchived data will be held. Verify and note the location for the Data Directory, and click **Next**. The *Choose Destination Location* screen for the Fusion PACS appears. (If the data directory folder does not exist, you will first be prompted to create it.)

Important: Make a note of the Data Directory location. The location of the Data Directory must be available when the server is set up in the Web GUI.

17. This is the location where the base server program will reside. All log and service files will also be in a subdirectory of this location. Verify and note the location for the Fusion PACS, and click **Next**. The *Local IP and AE Title Question* screen appears.
18. This information is used by the Web for remote access to the server through workstations and clients on the network. Enter the IP address and the AE Title of the server, and click **Next**. The *Local MSSQL Server Information* screen appears.

Important: Make a note of the local IP address and the DICOM name of the server. The user needs these values when it is time to configure the local system and remote workstations through the Web GUI.

19. Enter the IP address and the name of the SQL Server that Fusion PACS will use to store its database, and click **Next**. The *Check Setup Information* screen appears.
20. Verify that the installation information is correct, and click **Next**. The progress indicator indicates the progress of the installation process.

Once the main files have been copied, the setup of the remaining components begins.

Step 3: Completing the installation

Once you have installed the software, you are prompted to install the other software options for your selected installation type. The final steps are:

1. Install Web services and set up access to Key Images (see "Installing FUSION IIS Web services" on page 12).
2. Install NT services and configure a system user (see "Installing Windows NT Core services" on page 14).
3. Install or update the FUSION databases and select a date format (see "Installing the database" on page 17 and "Setting the date format" on page 18).
4. Select whether to secure the ImageChannel connection (see "Securing the Image Channel connection" on page 18).
5. Choose a workflow configuration and set up report access (see "Setting the FUSION Workflow configuration" on page 19).

Note: Steps 1, 2, and 3 only appear for selected installation types, as described in the first section.

Installing FUSION IIS Web services

This creates the virtual directory for the Web services in IIS. This step applies to complete installations and Web installations. This section describes the available options for installing the Web services.

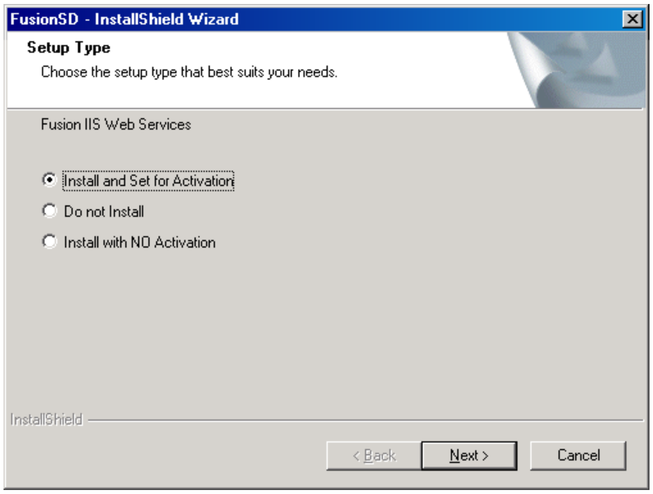

FusionSD - InstallShield Wizard

Setup Type

Choose the setup type that best suits your needs.

Fusion IIS Web Services

● Install and Set for Activation

○ Do not Install

○ Install with NO Activation

To install the Web services:

1. Select the type of setup for your installation, and click **Next**.

Install and Set for Activation	This selection must be used for all new installations. All of the web page components are installed, and the virtual directory in Internet Information Services (IIS) is created and run under **Default Website > FUSIONServer**.
Do not Install	No web page components are installed.
Install with NO Activation	All of the web page components are installed, but the virtual directory in ISS is not created.

Note: Install and Set for Activation must be enabled for a new installation. This service must be run on any new installations. Depending on the speed of the system, this may take some time. Please be patient and do not turn off your system during this process. Failure to enable **Install and Set for Activation** on new installations will result in failure of the system to create the virtual directory in Internet Information Services (IIS).

The installer proceeds to install and configure the Web services.

2. Once the Web services have been installed and configured, you will be asked whether you will be using key images.

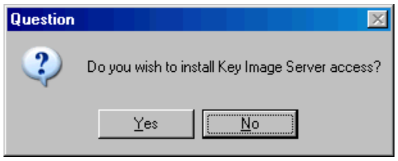

Question

Do you wish to install Key Image Server access?

3. If your site will be using key images, click **Yes**. You will asked to enter the URL for the Key Images server.

Note: To use Key Images, the eFilm Visualization Services must be installed on an accessible web server.

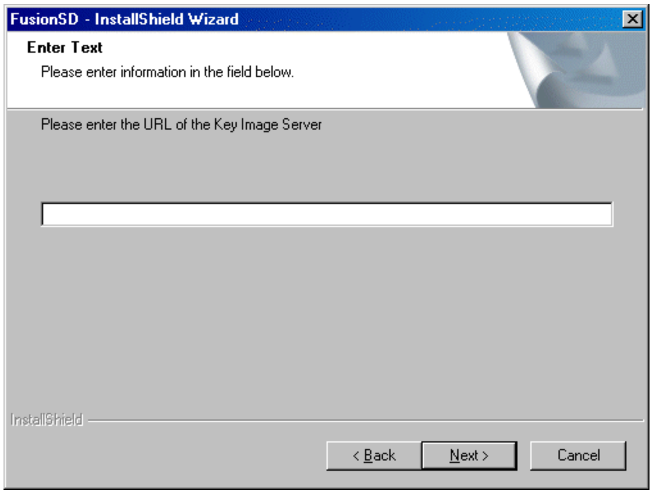

FusionSD - InstallShield Wizard

Enter Text

Please enter information in the field below.

Please enter the URL of the Key Image Server

4. Enter the URL to access the Key Images server. Typically, the URL for a Key Images server is:

```
http://<ipaddress>/VisualizationServices/KeyImages.asmx?wsdl
```

Note: You can verify the URL by typing it in a Web browser. You should receive an XML document in response.

5. Click **Next** to complete the Key Images setup.

Note: Before users can view key images, they must enable key images. See the *Storage and Distribution Manager User Guide* for information on enabling key images for users.

Installing Windows NT Core services

After you have installed the FUSION IIS Web services, you are prompted to install the FUSION Server Core services for Windows NT. This step applies to complete installations and Services only installations. This section describes the installation options.

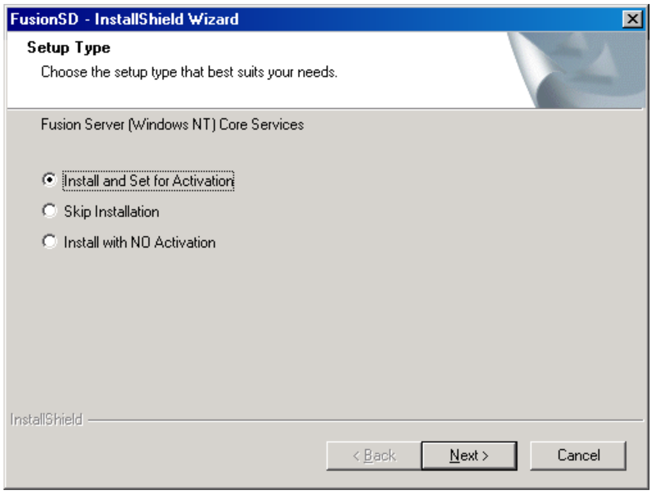

FusionSD - InstallShield Wizard

Setup Type

Choose the setup type that best suits your needs.

Fusion Server (Windows NT) Core Services

● Install and Set for Activation
○ Skip Installation
○ Install with NO Activation

To install the Windows NT Core services:

- Select the type of setup for your installation, and click **Next**.

Install and Set for Activation	This selection must be used for all new installations. All of the services are installed and set to active.
Skip Installation	No services are installed.
Install with NO Activation	All of the services are installed but are **not** set to active.

Important: Install and Set for Activation must be enabled for a new installation. This service must be run on any new installations. Depending on the speed of the system, this may take some time. Please be patient and do not turn off your system during this process.

The installer proceeds to install and configure the Windows NT Core services. Once the Window NT Core services have been installed, the system user account setup begins.

Setting up the system user

This section describes how to set up the system user account. This is the user account that the system will use. You should create the same user account on each server that will be running the FUSION services. The account should have the following properties:

- administrator privileges
- the ability to log in as a service

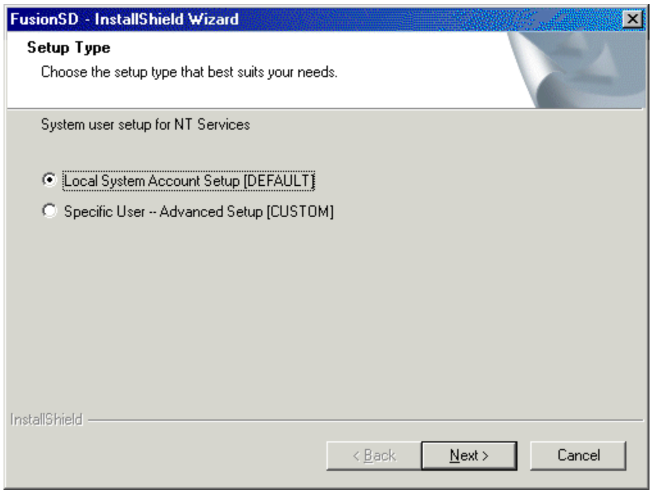
FusionSD - InstallShield Wizard
Setup Type
Choose the setup type that best suits your needs.
System user setup for NT Services
● Local System Account Setup [DEFAULT]
○ Specific User -- Advanced Setup [CUSTOM]
InstallShield
< Back
Next >
Cancel

Important: You must select **Specific User -- Advanced Setup (CUSTOM)** when using a network-attached storage (NAS) device.

To set up the system user:

1. Select the type of setup for your installation, and click **Next**.

Local System Account Setup (DEFAULT)	All services are registered under the local system account. Use this selection when all Fusion PACS components are being loaded on to the same server. For example, database, web page, services, and archiving all reside on the same server.
Specific User -- Advanced Setup (CUSTOM)	Allows you to select a common user for several machines if the Fusion PACS components have been divided between two or more servers; for example, if the database and web page do not reside on the same machine. **Note:** If required, make sure the user has administrator rights on both servers (Administrator account suggested--see above for more details) and make sure the password is the same for both servers.

   - If you selected **Local System Account Setup (DEFAULT)**, proceed with "Installing the database" on page 17.
   - If you selected **Specific User -- Advanced Setup (CUSTOM)**, continue to step 2.

2. A prompt appears, asking you to verify that the user exists as a member of the Administrators Group. Click **Yes**. The *Setting up FUSION Service User* screen appears.

If the user does not yet exist, you can leave the installer running while you create a suitable user account.

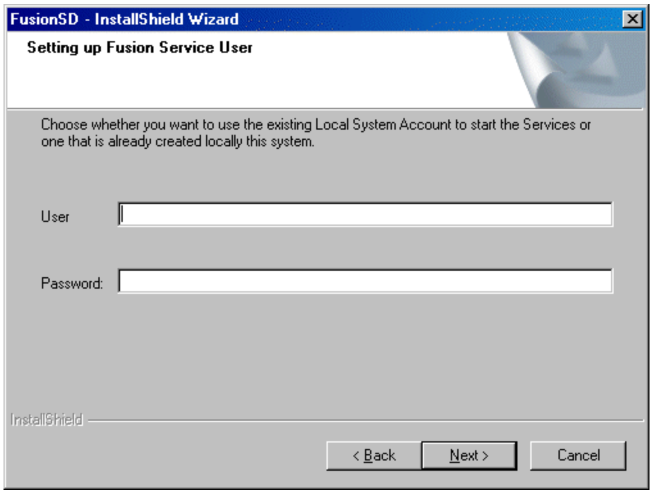

FusionSD - InstallShield Wizard

Setting up Fusion Service User

Choose whether you want to use the existing Local System Account to start the Services or one that is already created locally this system.

User

Password:

3. Enter a user with administrator rights on both servers and ensure that the password is the same for both servers. Click **Next**.

The system configures the user account for Windows NT services.

Installing the database

This section provides the procedure for installing the Fusion PACS database. This step applies to complete and database-only installations.

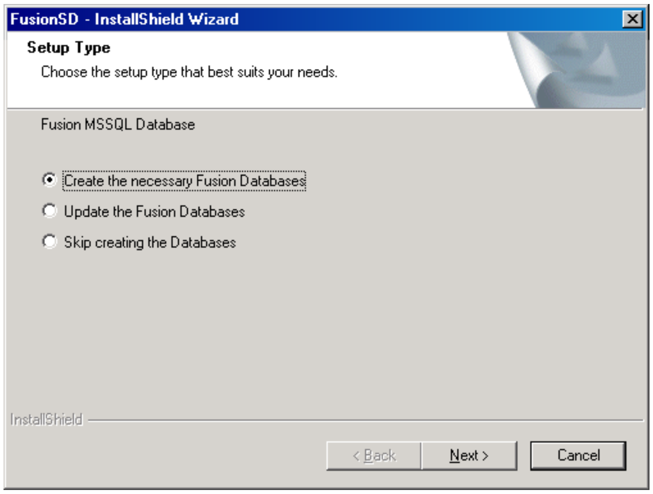

FusionSD - InstallShield Wizard

Setup Type

Choose the setup type that best suits your needs.

Fusion MSSQL Database

● Create the necessary Fusion Databases

○ Update the Fusion Databases

○ Skip creating the Databases

Important: The **Create the necessary FUSION Databases** selection will erase any previous FUSION database installation on the system.

To install the Fusion PACS Database:

4. Select the type of setup for your installation, and click **Next**.

Create the necessary FUSION Databases	Should be used for **new** installations. Installs new databases and creates all tables, stored procedures, and functions for both the Fusion PACS and Fusion PACS User Manager databases. It also inputs default values into specific tables.
Update the FUSION Databases	Should be used for **upgrades** from Fusion PACS version 1.4.2. Upgrades all tables, stored procedures, and functions, based on any changes that have been made.
Skip creating the Databases	Skips the database creation process. This should be used when installing the software for a clinical site.

Important: Create the necessary Fusion PACS Databases must be enabled for a new installation. This service must be run on any new installations. Depending on the speed of the system, this may take some time. Please be patient and do not turn off your system during this process. Failure to enable **Create the necessary Fusion PACS Databases** on new installations will result in failure of the system to create the selected SQL files that allow future upgrades to be performed.

The system reveals a series of status messages as it steps through the database installation process. You will then be prompted to select a date format.

Setting the date format

Select the format in which Fusion PACS will display dates. You can select either American (MM/DD/YYYY) or British (DD/MM/YYYY) format. Click a button to make your selection.

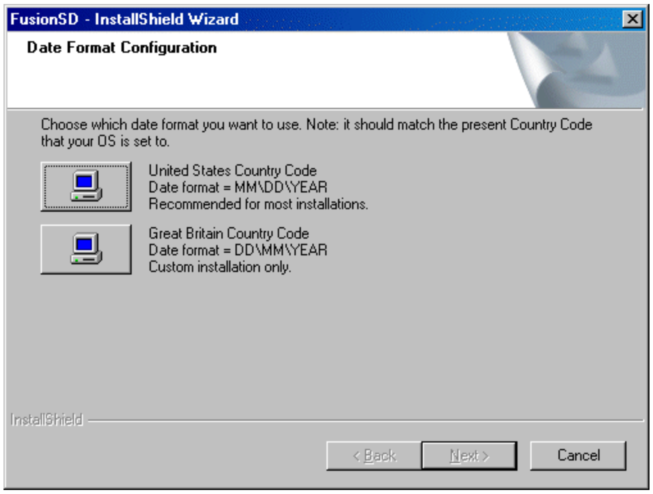

FusionSD - InstallShield Wizard

Date Format Configuration

Choose which date format you want to use. Note: it should match the present Country Code that your OS is set to.

United States Country Code
Date format = MM\DD\YEAR
Recommended for most installations.

Great Britain Country Code
Date format = DD\MM\YEAR
Custom installation only.

Securing the Image Channel connection

Once you have selected a date format, the *Secure Image Channel Connection* screen appears. This screen allows you to set up a secure connection to the Image Channel server using SSL instead of a VPN.

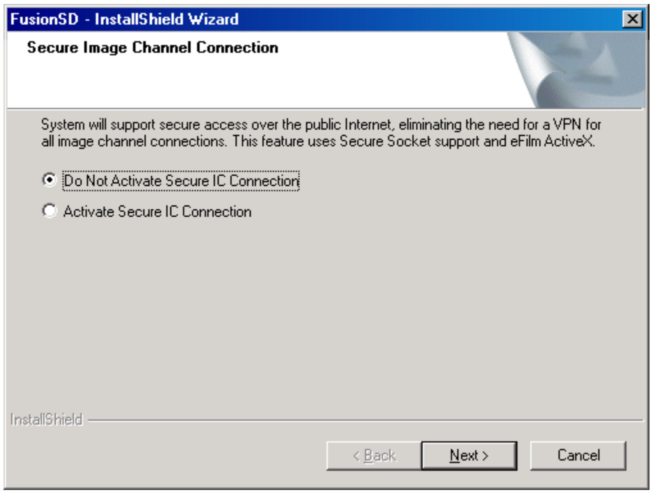

Secure Image Channel Connection

System will support secure access over the public Internet, eliminating the need for a VPN for all image channel connections. This feature uses Secure Socket support and eFilm ActiveX.

● Do Not Activate Secure IC Connection

○ Activate Secure IC Connection

To install a Secure Image Channel connection:

- Select the secure access option for your installation, and click **Next**.

Do Not Activate Secure IC Connection	The system will not activate a secure image channel connection.
Activate Secure IC Connection	The system activates a secure image channel connection.

The system activates the Image Channel, and indicates whether the connection is secure.

Setting the FUSION Workflow configuration

The final stage of installing Fusion PACS involves setting up the Fusion PACS Workflow configuration.

Note: Workflow and report integration settings have been added to the **web.config** file, which is set automatically during installation.

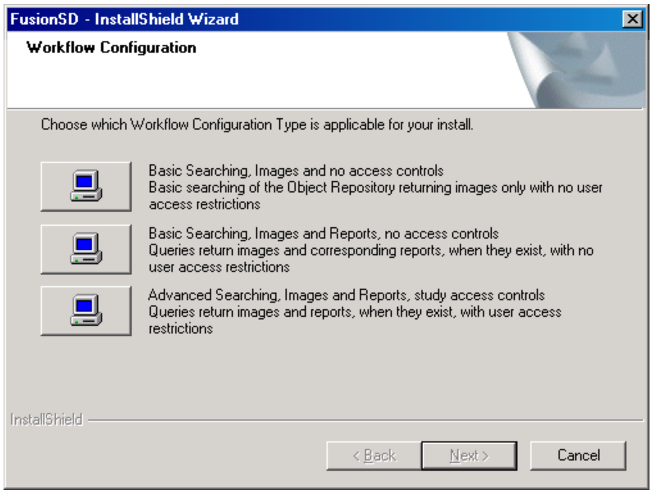

Workflow Configuration

Choose which Workflow Configuration Type is applicable for your install.

Basic Searching, Images and no access controls
Basic searching of the Object Repository returning images only with no user access restrictions

Basic Searching, Images and Reports, no access controls
Queries return images and corresponding reports, when they exist, with no user access restrictions

Advanced Searching, Images and Reports, study access controls
Queries return images and reports, when they exist, with user access restrictions

To set up the FUSION Workflow Configuration:

1. Select the type of setup for your installation

Basic Searching, Images and no access controls	Installs functionality for basic searching of the PACS, returning images only with no user access restrictions.
Basic Searching, Images and Reports, no access controls	Fusion PACS Workflow allows you to query for reports at a URL. You must configure the HTTP service with an appropriate IP address and port number.
Advanced Searching, Images and Reports, study access controls	Fusion PACS Workflow provides user authentication and allows users to view reports according to their access privileges. In addition to configuring the HTTP service with an appropriate IP address and port number, you must configure the Web Query Retrieve service on the Workflow server.

2. If you selected the first option, you're done! Continue the configuration in Chapter 2.

   If you selected the second option, go to step 4 (Report Server setup).

   If you selected the third option, go to step 3 (Workflow Server setup).

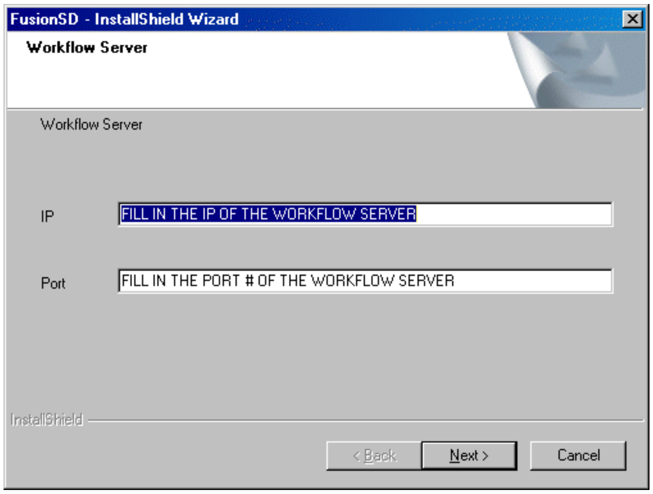

Workflow Server

IP FILL IN THE IP OF THE WORKFLOW SERVER

Port FILL IN THE PORT # OF THE WORKFLOW SERVER

3. Enter the IP address and port number of the workflow server and click **Next**. The *Report Server* screen appears.

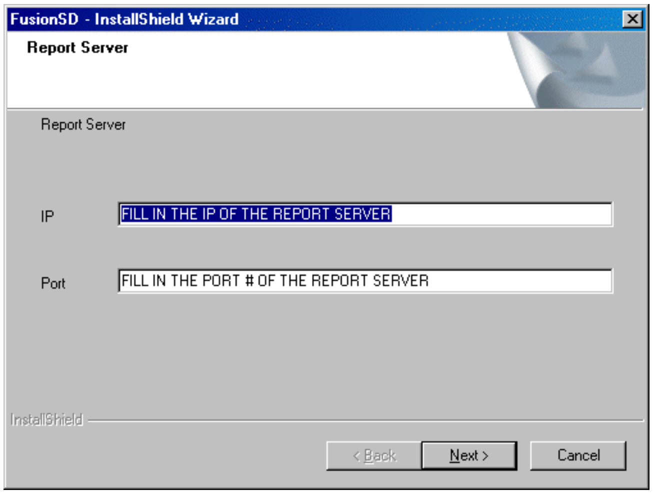

FusionSD - InstallShield Wizard

Report Server

Report Server

IP: FILL IN THE IP OF THE REPORT SERVER

Port: FILL IN THE PORT # OF THE REPORT SERVER

4. Enter the IP address and port number of the report server and click **Next**. The Web URL screen appears.

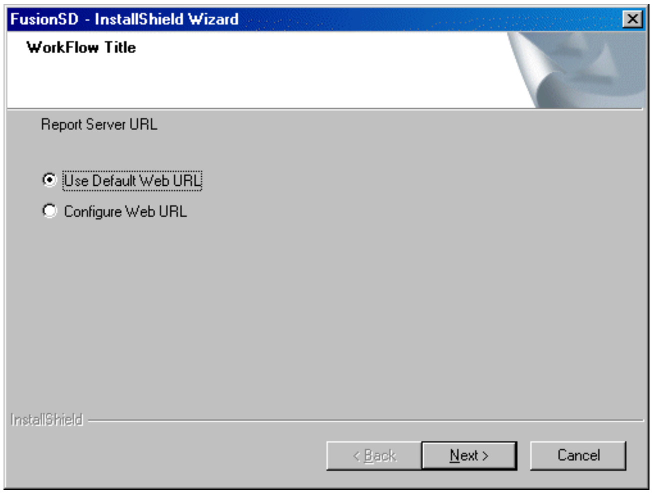

FusionSD - InstallShield Wizard

WorkFlow Title

Report Server URL

● Use Default Web URL

○ Configure Web URL

5. Select whether to use the default URL to submit reports or enter a custom URL. If you use the default, you're done! Otherwise, the report URL screen appears.

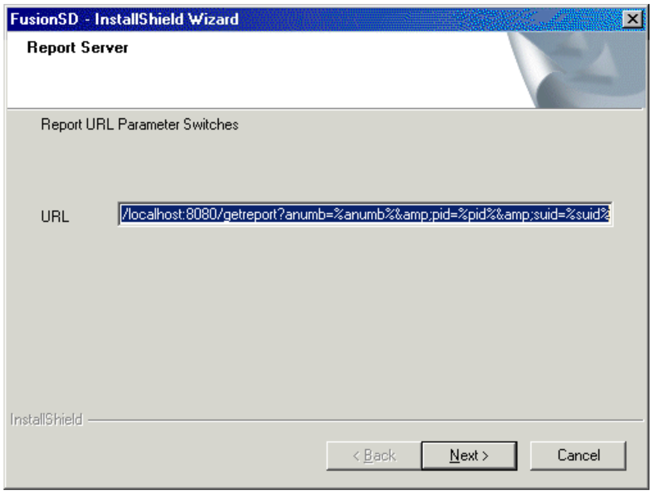

FusionSD - InstallShield Wizard

Report Server

Report URL Parameter Switches

URL /localhost:8080/getreport?anumb=%anumb%&amp;pid=%pid%&amp;suid=%suid%

6. Enter the URL on which to submit report requests and click **Next**.
7. You're done! Click **Finish** to exit the wizard. The Fusion PACS installation is now complete.

Configuring Fusion PACS 2

This chapter provides instructions for configuring special options for Fusion PACS. It shows you how to:

- complete the basic configuration steps required to use the PACS (see “Basic PACS configuration” on page 23)
- determine which move page the PACS will use (see “Setting the move page” on page 23
- enable n-tier file storage, if purchased (see “Enabling n-tier storage” on page 23)
- set up and operate SQL log shipping (see “Using SQL log shipping” on page 24).
- secure the PACS (see “Securing Fusion PACS” on page 35).
- set up load balancing (see “Setting up load balancing” on page 41).

Basic PACS configuration

Following the instructions in the *Fusion PACS Storage and Distribution Manager User Guide*, you must set up, at a minimum, a file system and a server to represent the installed Fusion PACS machine in the Web GUI.

Setting the move page

Two different move pages are available for use with Fusion PACS. You can choose which page to use by editing the web.config file.

Note: This step is **not** part of the installer, and must be performed manually on the server.

To change the move page setting:

1. On the machine where Fusion PACS is installed, locate the **web.config** file in the directory where you installed Fusion PACS and open it using a text editor such as Notepad.
2. Edit the following parameter, setting the value to **true** to use the newer move page:

   ```
   <add key=”ActivateNewMovePage” value=”true”/>
   ```
3. Save and close the file.

Enabling n-tier storage

This step is only necessary if the customer has purchased the n-tier storage functionality.

To enable n-tier storage:

1. Open *Query Analyzer* and select the **FusionServer** database.
2. Run the **n-Tier_Enable.sql** script.

Important: The n-Tier_Enable.sql script is not provided with the installed software. You must obtain the script from Merge Healthcare Service.

Using SQL log shipping

The critical flaw with single-server implementations of Fusion PACS is that if anything happens to that server, even if it's been backed up regularly, you're going to lose a lot of data and a fair amount of time getting a replacement box in place and configured. SQL Server Log Shipping allows us to create a hot database backup on a separate SQL Server in case the production database server fails for any reason.

SQL log shipping works like this:

1. SQL Server backs up the database and transaction logs to a folder on the production server.
2. A SQL Agent job running on the production server copies the backed-up data to the backup server.
3. (Optional, if the restore function is enabled) The backup SQL Server restores the database from the folder on the backup server.

In the event that the production database server fails for any reason, the backup database can be promoted to the primary role with minimal disruption.

This section describes:

- how to install SQL log shipping (see "Installing SQL log shipping" on page 24)
- how to configure SQL log shipping (see "Configuring SQL log shipping" on page 28)
- how to monitor the status of SQL log shipping (see "Monitoring SQL log shipping" on page 30)
- how to switch to the backup database (see "Switching to the backup database" on page 31)
- how to set up a new backup database (see "Setting up a new backup server" on page 33)
- what to do if the transaction log backup/restore process fails (see "Transaction Log Backup Failure and Recovery" on page 34)
- how to synchronize the production and backup databases (see "Synchronizing the databases" on page 34)
- how to uninstall log shipping (see "Uninstalling log shipping" on page 35)

Installing SQL log shipping

SQL log shipping essentially operates as a series of SQL Agent jobs that handle backing up the production database, copying the backups, and restoring backups to the secondary database server. These jobs are installed by configuring and running two SQL scripts.

Note: This procedure assumes that you have already set up two database servers, and that one of those database servers has been set up as the Fusion PACS database server. Ideally, the other components of Fusion PACS are on a third server, entirely separate from the database servers.

To install SQL log shipping:

1. Create two folders:
   - On the **production database server**, create a folder in which to store the database backups. This folder needs to be accessible and given write permissions for the user account that is running the production SQL Server Agent service.
   - On the **backup database server**, create the folder to which the database backups will be copied. This folder must be shared and given the security rights so the production SQL Server Agent service user account can write to this folder.

   Important: Ensure that there are **no spaces** in the path to either folder. The SQL Agent job will fail if the log shipping process attempts to copy a file from or to a path with spaces.
2. On the production database server, configure and run the two SQL scripts that will create the necessary SQL Agent jobs (see "Configuring the SQL scripts" on page 25).
3. Run the SQL Agent jobs (see "Running the SQL Agent jobs" on page 28).

Configuring the SQL scripts

Two SQL scripts must be executed on the FusionServer production SQL Server master database. It is not necessary to execute the scripts on the backup SQL server; the scripts will create the stored procedures used for the backup/restore process on both servers at the same time.

The SQL Agent jobs and procedures being installed are designed to work in one of two ways. The token defined as "Controls if the restore process occurs" needs to be set accordingly to the type of SQL Log Shipping install you wish to deploy:

- To maintain a hot backup SQL server, set the value of the token to 1 (one).

  Note: With this type of installation, all tokens in the script file must have valid values.
- To perform database and transaction log backups only (i.e., you are not running a backup server), set the value of this token to 0 (zero). With this type of install, not all tokens will need to be replaced. The definition of each token indicates whether it is needed when the restore process is off.

  In most instances of the log shipping process, a backup server should exist and the restore process should be turned on. But if you're down to one server for whatever reason, you'll still want to keep making the backups until you get your second machine online again. In that case, the log shipping scripts can be executed on the new production server with the restore process turned off. This allows some means of backup until the entire system can be stabilized, and allows for easier restore of the old production server when it is able to resume the role of main server. This is explained more in "Switching to the backup database" on page 31.

LogShipping_FusionServer.sql

This script installs two SQL Agent jobs:

- The **Fusion Server Database Backup** job backs up the database, copies the backup file to the specified folder, and (if restore is on) restores the database on the backup server. By default, this job runs each night at 1:00 am.
- The **Fusion Server Transaction Log Backup** job backs up the FusionServer transaction log, copies the backup file to the specified folder, and (if restore is on) restores the transaction log onto the backup database. By default, this job runs every ten (10) minutes.

The following table lists the Key Tokens that must be replaced before executing the script:

Token	Description
AAAAA	The IP address or hostname of the production SQL server.
BBBBB	The name of the production FusionServer database.
CCCCC	The user name of the database administrator (sa) account on the Production Server.
DDDDD	The password of the database administrator (sa) account on the Production Server.
EEEEE	The path to the Production Server backup folder.
FFFFF	The IP address or hostname of the backup SQL Server.*
GGGGG	The name of the backup FusionServer database.*
HHHHH	The user name of the database administrator (sa) account on the Backup Server.
IIIII	The password of the database administrator (sa) account on the Backup Server.
JJJJJ	The location of the Backup Server Storage Folder From Production.* This should be in the format `<IP address>/<share name>`
KKKKK	The path to the folder on the Backup Server where the database backups will be copied.*
LLLLL	BackupServerSQLDataFile1*
MMMMM	BackupServerSQLDataFile2*
NNNNN	BackupServerSQLDataFile3*
OOOOO	BackupServerSQLLogFile*
PPPPP	BackupServerLogicalDataName1*
QQQQQ	BackupServerLogicalDataName2*
RRRRR	BackupServerLogicalDataName3*
SSSSS	BackupServerLogicalLogName*

Token	Description
TTTTT	Flag if restore process is turned on (1 or 0)
UUUUU	This is the number of days to keep database backups. Normally, this value should be set to 5 for one complete workweek. The FusionServer database can become quite large, so it is necessary to set this token to a reasonable value as to not drain disk space on the backup drives.

*Not required if restore is off.

LogShipping_FusionServerUserManager.sql

This script installs two SQL Agent jobs:

- The **Fusion Server User Manager Database Backup** job backs up the database, copies the backup file to the specified folder, and (if restore is on) restores the database on the backup server. By default, this job runs each night at 4:00 am.
- The **Fusion Server User Manager Transaction Log Backup** job backs up the FusionServer transaction log, copies the backup file to the specified folder, and (if restore is on) restores the transaction log onto the backup database. By default, this job runs every twenty (20) minutes.

The following table lists the Key Tokens that must be replaced before executing the script:

Token	Description
AAAAA	The IP address or hostname of the production SQL server.
BBBBB	The name of the production FusionServerUserManager database.
CCCCC	The user name of the database administrator (sa) account on the Production Server.
DDDDD	The password of the database administrator (sa) account on the Production Server.
EEEEE	The path to the Production Server backup folder.
FFFFF	The IP address or hostname of the backup SQL Server.*
GGGGG	The name of the backup FusionServerUserManager database.*
HHHHH	The user name of the database administrator (sa) account on the Backup Server.
IIIII	The password of the database administrator (sa) account on the Backup Server.
JJJJJ	The location of the Backup Server Storage Folder From Production.* This should be in the format `<IP address>/<share name>`

Token	Description
KKKKK	The path to the folder on the Backup Server where the database backups will be copied.*
LLLLL	BackupServerSQLDataFile*
MMMMM	BackupServerSQLLogFile*
NNNNN	BackupServerLogicalDataName*
OOOOO	BackupServerLogicalLogName*
PPPPP	Flag if restore process is turned on (1 or 0)
QQQQQ	This is the number of days to keep database backups. Normally, this value should be set to 5 for one complete workweek. The FusionServer database can become quite large, so it is necessary to set this token to a reasonable value as to not drain disk space on the backup drives.

Running the SQL Agent jobs

After successfully installing these scripts, you should then run the newly installed jobs:

1. In SQL Server Enterprise Manager navigate to the **Production SQL Server > Management > SQL Server Agent >** highlight the **Jobs** node. You should see a list of the newly created Log Shipping jobs on the right pane.
2. Run both the **Fusion Server Database Backup** and **Fusion Server User Manager Database Backup** jobs. Running these jobs will prime the backup server databases. Failure to do so will result in the Transaction Log Backup jobs failing.

Note: After restoring the database backup to the backup server, the backup database status will be "Loading". This is because SQL Server requires that databases be non-operational when restoring multiple transaction logs. The backup database should remain non-operational until you need to promote it to production status. For information on how to change a "Loading" non-operational database to a normal operation database, see “Switching to the backup database” on page 31.

Configuring SQL log shipping

This section describes some of the post-installation configuration changes you may want to make to the SQL Agent jobs. It describes how to:

- change when the jobs run (see “Changing job schedules” on page 28)
- set up alerts to notify you of problems or events (see “Setting up alerts” on page 29)

Changing job schedules

All SQL Agent jobs are installed with a default schedule. These schedules may be modified to suit the needs of the site.

To modify the schedule for a job:

1. Open SQL Server Enterprise Manager and connect to the production database.

2. Navigate to **Management\SQL Server Agent\Jobs**.
3. Right-click the desired job and select **Properties**.
4. Select the **Schedules** tab. There will be one schedule item for each job. Select this item and click **Edit**. The job schedule edit window will open where you can modify the schedule for the selected job.

Scheduling tips

Scheduling database backups:

- Database backup jobs should run once or twice a day; often enough to ensure that you have full recent backups of your database, but no so often that you affect system performance.
- Database backups should be scheduled for times when the SQL Server does not have a lot of user connections or activity.

Scheduling transaction log backups:

- Transaction log backup jobs should run every 10 to 15 minutes.
- Because the FusionServer database will generally be busier than the FusionServerUsermanager database, the FusionServer Transaction Log Backup job is (by default) scheduled to run every 10 minutes, and the FusionServerUserManager Transaction Log Backup job runs every 20 minutes.

Setting up alerts

We highly recommend that an alert(s) be set up for each job in order to monitor the status of each job.

There are two types of alerts that can be configured: alerts and notifications. Alerts will allow you to detail certain events when an error or certain event happens and execute a response for the event. Notifications allow you to configure a notification service to a pre-configured SQL Server Operator when the job hits a certain status during execution.

Both of the alert types are useful and should be used to create a complete notification service to help monitor the status of the log shipping jobs; for example, you could set up an alert for the FUSION Server Database Backup job that will alert a SQL Server operator if the job has encountered a specific error number. A custom text message can be emailed, sent to a pager, or a network messaged to the operator.

The FUSION Server Database Backup job can also have notifications assigned that will notify SQL Server operators when the job completes, fails, or succeeds. This can ensure that all jobs are running appropriately, or if they fail for some reason.

For more information on the proper way to configure and set up Alerts and Notifications, please refer to SQL Server Books Online.

To set up alerts for a job:

1. Open SQL Server Enterprise Manager and connect to the production database.
2. Navigate to **Management\SQL Server Agent\Jobs**.
3. Right-click the desired job and select **Properties**.
4. To create an alert, select the **Schedules** tab and click **New Alert**.
5. To create a notification, select the **Notifications** tab.

6. Click **OK** when you have made your changes.

Monitoring SQL log shipping

This section describes the tools you can use to monitor the status of the log shipping jobs:

The most important way to monitor the status of the log shipping jobs is to set up Alerts and/or Notifications of the execution of the job. This will provide as close to real time status information to be delivered.

If the log shipping job for the FusionServer transaction log has not been running successfully for 4 hours and the production server fails, you could potentially have lost 4 hours of data. With the proper Alerts and/or Notifications setup, some automatic action could take place the first time the transaction log backup job fails. This will allow the database administrator to find out what the problem is and to resolve it before a catastrophe happens.

SQL Server Enterprise

You can monitor the status of each log-shipping job inside SQL Server Enterprise Manager. If you navigate to the Jobs folder you will see a list of the SQL Agent jobs currently installed, including the log shipping jobs. If a job has been running successfully, you will see a normal window icon next to the job. If the job has failed for some reason, you will see a Red X next to the job.

You can find more detailed information on the history of the execution of a job by navigating to the job, highlighting the job, right mouse clicking on the job, and selecting View Job History. This will display the history of the execution of the job. Each job should have a Result of Successful.

Alerts/Notifications

If the log-shipping job has been configured with an alert, or notification of an action, then some kind of predefined message will be delivered to notify the receiver of the job status. This log shipping job status can be that the job was successful, the job failed, or if the job failed with a particular type of error.

SQL Server Query Analyzer

You can obtain status information on a log shipping job by executing a Transact SQL statement to return a result set giving information about the job. Running the following statement will return the execution history of the job.

```
EXEC Msdb.dbo.sp_help_jobhistory 'NameOfTheJob'
```

SQL Server Logs

The SQL Server logs on the production server and the backup server will show detailed information on the status of the backup/restore process of log shipping. You can coordinate the data from the log shipping job history with the SQL Server Log data to make sure the most current transaction and database backups have been restored on the backup database.

Windows Event Viewer Application Logs

By default, the log shipping jobs are configured to output status information to the Windows Event Viewer Application Log.

Switching to the backup database

This section describes the procedure to follow when the production database or server is unable to process requests for data. One assumption must be made while reading this section: that the backup server is in synch with the production server. All backup log-shipping jobs have completed successfully and the backup database is a mirror copy of the production database up to the last transaction log backup and restore.

If the backup server is out of synch with the production server because of failed log shipping jobs, you must re-synchronize the databases (see "Synchronizing the databases" on page 34).

When the production database has failed, the database administrator must quickly identify if the production server problems can be resolved, or if the backup SQL Server must be promoted to the new production server. Use the following flowchart to determine what you must do:

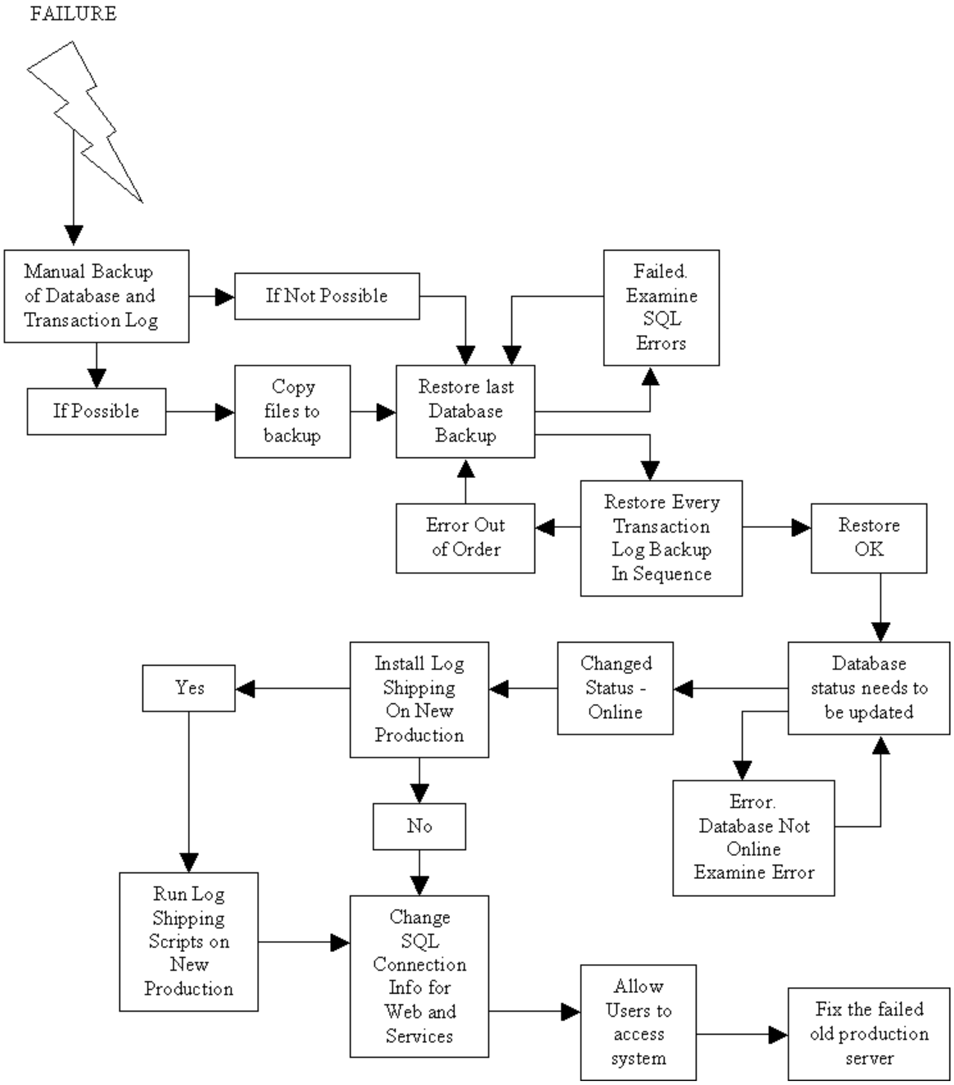
FAILURE
Manual Backup of Database and Transaction Log
If Not Possible
Failed. Examine SQL Errors
If Possible
Copy files to backup
Restore last Database Backup
Error Out of Order
Restore Every Transaction Log Backup In Sequence
Restore OK
Yes
Install Log Shipping On New Production
Changed Status - Online
Database status needs to be updated
No
Error. Database Not Online Examine Error
Run Log Shipping Scripts on New Production
Change SQL Connection Info for Web and Services
Allow Users to access system
Fix the failed old production server

If the production server has been fixed and is running normally, a quick verification of the log shipping jobs must be executed to ensure the backup server is being kept in synch with the production server.

If the production server is down for a significant amount of time, the backup server will need to be promoted.

To fail over to the backup server:

1. The first thing to attempt to do is a manual backup of the transaction log on the production server. This includes both the FusionServer and FusionServerUserManager databases. This can only be accomplished if the SQL Server is still running and somewhat accessible. The transaction log will need to be copied to the backup server, so network connectivity is required, unless a portable storage device is available. If the manual transaction log backup on the production server is successful, it can be restored on the backup server. This will ensure that the failed production server and backup server are mirror images of each other.
2. If the production server transaction logs cannot be backed up and restored to the backup server, then there will be a period of time where data will be lost. The amount

of time depends on the period of time from the last transaction log backup/restore and the production server failure.

3. The backup databases status will need to be modified from "Loading" and non operational to an operational state. This will allow users to access the databases and resume their work. To modify the status of the databases on the backup server, execute the two Transact SQL statements below on the backup SQL Server master database.

   ```
   RESTORE DATABASE FusionServer WITH RECOVERY
   RESTORE DATABASE FusionServerUserManager WITH RECOVERY
   ```

   For more information regarding this step, please refer to the Microsoft article located at http://support.microsoft.com/?kbid=822852

4. Once the databases on the backup server are operational a decision needs to be made if the new production databases and transaction logs will need to be backed up and stored on another physical storage device. This is highly recommended since the backup databases could fail at any minute leaving hours, or days of data lost. This will also provide a way to restore the old production server once it is available.

   "Installing SQL log shipping" on page 24 describes the log shipping installation process. Please refer to this section and take special note to options for backing up database and transaction logs without the restore option.

5. When the backup databases are operational and as up to date as they can be, then clients can begin to access the new production server. Connection information for all clients accessing data will need to be modified to connect to the new production server. This can be as easy as an IP address change, or a name change in a configuration file.

6. Once clients are accessing the new production server, it is vital that the failed server be fixed and brought back on line as soon as possible (see "Setting up a new backup server" on page 33). The system is in a very volatile state since there is no backup server to fail to. If the backup server fails for some reason, all data access for clients will cease to exist.

Setting up a new backup server

If a failover occurs it will be very important to get the old production server fixed, or another production server staged in a timely fashion. Once this is done, the old production server and databases will need to be recovered to mirror the current production server.

To set up a new backup server:

1. If the log shipping jobs have been executing on the current production server, then database backups and transaction logs have been occurring. If log shipping was not installed on the current production server, then manual database backups and transaction logs will need to be created.
2. Find, or create the last good full database backups on the current production server for the FusionServer and FusionServerUserManager databases. Restore these backups onto the old production server.
3. Begin to restore every transaction log backup for each FusionServer database since the point in time when the last full database backup was executed in the same order the transaction logs were created. If the transaction log restoring process gets out of synch, the entire process will need to be re started.

4. While all of this is going on the current production server is still processing transactions. At some point all users will need to be notified that they need to get out of the system. When all users are out of the system create a final transaction log on the current production server for each FusionServer database and restore them to the old production server you are staging to bring on line.
5. The current production server will need to be demoted to the backup role. In order to do this, if log shipping was installed on this server the current log shipping jobs will need to be uninstalled from the server. See "Uninstalling log shipping" on page 35 for information on how to remove log shipping.
6. The server being staged to take the role of the production server will need to have log shipping installed on it. See "Installing SQL log shipping" on page 24.
7. The new production server is now ready to be brought back online. All client connection information will need to be modified to connect to the new production server.
8. Once the new production server is operational verify that the log shipping jobs are executing successfully. If they are, then the production server and backup server are in synch and the system is back to normal.

Synchronizing the databases

If the log shipping process is failing for some reason and the backup databases are out of synch with the production databases you will need to re synch the two. This process will ensure that the production databases and backup databases are mirrors of each other.

To synchronize the databases:

1. Find the last full database backup. This can be for the FusionServer, or FusionServerUserManager databases if they are both out of synch.
2. Restore these backups on the backup SQL Server.
3. Restore every transaction log backup for each FusionServer database (if necessary) since the point in time when the last full database backup was executed in the same order the transaction logs were created. If the transaction log restoring process gets out of synch, the entire process will need to be re started
4. Once this is done, the production and backup databases should be in synch with each other.
5. Explore why the log shipping process was failing. Once this is resolved, verify that the log shipping process is functioning properly and keeping the two servers in synch with each other.

Transaction Log Backup Failure and Recovery

In the event that the transaction log backup and restore job for both databases has failed; the stored procedure that executes these tasks can be invoked manually with a parameter to attempt to restore the missing transaction logs.

By running:

```
EXEC spLS_FusionServerTransactionLogBackupJob 1
```

Will attempt to automatically resynchronize the backup FusionServer database with the missing transaction logs

```
EXEC spLS_FusionServerUserManagerTransactionLogBackupJob 1
```

Will attempt to automatically resynchronize the backup FusionServerUserManager database with the missing transaction logs

The parameter sent in is a flag that allows the stored procedure to examine the backup history of the production SQL Server and match it with the restore history on the backup server. The missing transaction logs on the backup server will go through an attempt to automatically be restored in the proper sequential order.

In order for this to occur, the transaction log backup files need to exist on the restore file share where the backup files are copied. If for some reason they are not present, then the automatic restore will fail.

Uninstalling log shipping

If there is a need to uninstall the log shipping files and procedures you will need to run the script file **LogShipping_FusionSever_Uninstall.sql** on the production FusionServer SQL Server. This file will uninstall all stored procedures and SQL Agent jobs for the log shipping process on the FusionServer and FusionServerUserManager databases.

The uninstall script file has key tokens that will need to be replaced before execution of the script. The key tokens are defined in each script file, but are also defined under each script file definition below.

Key Token	Description
AAAAA	FusionServer Production SQL Server Name
BBBBB	Production SQL Server Administrator User Name
CCCCC	Production SQL Server Administrator password for user in BBBBB
DDDDD	Backup SQL Server name (not required if restore is off)
EEEEE	Notifies the uninstall process if the log shipping install has the restore backup file turned on. ▪ If the restore process is on, set this token to 1 (one). ▪ If the restore process has not been installed, set this token to 0 (zero).

Securing Fusion PACS

If Fusion PACS is running in conjunction with a VPN and no other type of external access is required, then no further steps need be taken for security. However, if the site requires HTTPS access to Fusion PACS, you can follow the instructions provided here to set up secure access.

To configure Fusion PACS for HTTPS access:

1. Install Fusion PACS and activate the secure Image Channel connection (see “Securing the Image Channel connection” on page 18).
2. Secure the Web page for external access (see “Securing the Web page” on page 36).
3. Secure the Image Channel connection (see “Securing SSL connections for Fusion PACS” on page 38).

4. Create a PACS-only database user and encrypt the database access information (see "Securing database access" on page 39).

Securing the Web page

There are three steps to securing the Web page for external access:

1. Configure IIS to place web communication on port **443 [https://])**.
2. Prepare a certificate request (see "Preparing a certificate request" on page 36).
3. Using the certificate request, obtain a certificate (see "Requesting and installing certificates" on page 37).

Preparing a certificate request

Before you can obtain a certificate, you must prepare a certificate request.

To prepare a certificate request:

1. In Windows, select **Start > Settings > Control Panel > Administrative Programs > Computer Management**. The *Computer Management* application starts.
2. Expand **Internet Information Services** (IIS)> **Default Web Site**.
3. Right-click and select **Properties**. The *Default Web Site Properties* screen appears.
4. Click the **Directory Security Tab** and select the **Server Certificate button**. The *Web Server Certificate Wizard* begins.

   The Web Server Certificate Wizard produces a certificate request, which is required to obtain a certificate from a certificate authority (CA). Your site may operate its own certificate authority using Microsoft Certificate Services (in which case, see the procedure below) or obtain its certificates from an independent CA such as VeriSign.
5. Click **Next** to continue. The *Server Certificate Method* screen appears.
6. Select **Create a new certificate**, and click **Next**. The *Delayed or Immediate Request* screen appears.
7. Select **Prepare the request now, but send it later**, and click **Next**. The *Name and Security Settings* screen appears.
8. Enter the security certificate name (e.g., **Merge eFilm**), and click **Next**. The *Organization Information* screen appears.
9. Enter the name of your organization and the name of the organization's division or department, and click **Next**. The *Your Site's Common Name* screen appears.
10. Enter the hostname of the Fusion PACS machine, and click **Next**. The *Geographical Information* screen appears.
11. Enter the institution's address information, and click **Next**. The *Certificate Request File Name* screen appears.
12. Enter a name for the file. This will create an output file that you can save in any accessible location. Click **Next**. The *Request File Summary* screen appears.

    Note: You may need to email a copy of the Certificate Request File to a certificate authority, if your institution is not running its own.
13. Click **Next**, then click **Finish** to close the *Web Server Certificate Wizard*.

Requesting and installing certificates

There are two options for requesting and installing certificates:

- If you are using a third party certificate authority (CA) such as Verisign, see "Using a third-party CA" on page 37.
- If your organization has its own CA, see "Using an internal CA" on page 37.

Using a third-party CA

If your institution does **not** run its own CA, follow the procedure below to issue and install the certificate.

To request a third-party certificate:

1. Email the certificate request created in the previous procedure ("Preparing a certificate request" on page 36) to an independent CA such as Verisign.
2. An email will inform you which web page to visit to download your certificate once you have been authenticated. This process will typically take several days.
3. Once you have downloaded the certificate, you are ready to proceed.

Using an internal CA

If your organization has its own CA, follow the procedure below to issue and install the certificate.

To request and install a Microsoft Certificate for the Web page:

1. Connect to the Intranet address for your CA. The *Welcome* screen appears.
2. Select **Request a certificate**. Click **Next**. The *Choose Request Type* screen appears.
3. Select **Advanced request**. Click **Next.** The *Advanced Certificate Requests* screen appears.
4. Select **Submit a certificate request using a base64 encoded PKCS #10 file** or a renewal request using a base64 encoded PKCS #7 file.

   Note: The selection above is used because a txt file was exported.

   Click **Next**. The *Submit a Saved Request* screen appears.
5. Select **Browse** and navigate to your saved text file (as created in the previous procedure), or open the text file and then copy and paste its contents into the supplied field.
6. Select **Submit**. The system authenticates your request and advances to the *Certificate Issued* screen.
7. Select your certificate file. The *File Download* screen appears.
8. Save the .cer file to the C drive (or system drive). The *Default Web Site Properties* screen appears.
9. Select **Server Certificate**. The *Web Server Certificate Wizard* starts.
10. Click **Next**. The *Pending Certificate Request* screen appears.
11. Select **Process the pending request and install the certificate**.

    Note: You can also delete pending certificate requests through this screen.

    Click **Next**. The *Process a Pending Request* screen appears.

12. Select **Browse** and navigate to the .cer file that you downloaded from your certificate authority. Click **Next**. The *Certificate Summary* screen appears.
13. Click **Next** to install the certificate on your server. The *Completing the Web Server Certificate Wizard* screen appears.
14. Click **Finish** to close the Web Server Certificate Wizard. The *Default Web Site Properties* screen appears.
15. Click **Edit** under Secure communications. The *Secure Communications* screen appears.
16. Select **Require secure channel (SSL)**. Click **OK**. Web connections on this server are secured.

Securing SSL connections for Fusion PACS

This section describes how to secure an SSL connection for the Image Channel server.

To secure an SSL connection for the Image Channel:

1. Connect to a Certificate Server as above ("Requesting and installing certificates" on page 37) and request a certificate.
2. Under advanced request type select **Submit a certificate request using a form**. The *Advanced Certificate Request* screen appears.
3. Fill out the form with your name, email, company, department, city, state, country, and select **Mark keys as exportable**.
4. A Potential Scripting Violation message appears. Click **Yes** to proceed with the certificate request. The *Certificate Issued* screen appears.
5. Click **Install this Certificate**. The system installs the certificate in Internet Explorer.
6. Open Internet Explorer and select **Tools**; **Internet Options**; **Content**. The *Internet Options* screen appears.
7. Click **Certificates** and click the **Personal** tab. The *Personal Certificates* screen appears.
8. Highlight your certificate (which should be called Merge eFilm) and click **Export**. The *Certificate Export Wizard* starts.
9. Click **Next**. The *Export Private Key* screen appears.
10. Click **Yes, export the private key**. Click **Next**. The *Export File Format* screen appears.
11. Click **Next**. Do not change the format. The *Password* screen appears.
12. Type in the password (default is **password**).

    Note: The password is a hardcoded value. To use a different password (for security purposes), you must edit the following registry key: **HKEY_LOCAL_MACHINE\Software\Merge eFilm\Fusion Server\CertificatePassword**

    The *File to Export* screen appears.
13. Type **fusion.pfx** in the **File name** field and browse to **C:\Program Files\Merge eFilm\Fusion Server**. The server automatically looks for this filename during installation.

    Note: To change the file name you must edit the following registry key: **HKEY_LOCAL_MACHINE\Software\Merge eFilm\Fusion Server\CertificateFile**
14. Click **Next**, then click **Finish**. The SSL IC is configured.

Note: In the case of a third party certificate authority, there may be vendor-specific actions required to export the private key. Consult with the vendor for more information.

To verify that the connection is secure:

1. Enter the Web GUI and view a study.
2. When the image is displayed, select the bottom portion of the image (a black area beside the toolbar) and select **View > Source** from your Web browser's menu bar.
3. The following lines in the file should have the values displayed below, and indicate the port number being used by Image Channel:

```
CFGMETADATASERVERPORT = 4444
CGFPIXELDATASERVERPORT = 4444
CFGPIXELDATASECURE = 1
CFGMETADATASECURE = 1
```

Securing database access

By default, Fusion PACS stores the database user ID and password in three places:

- the **web.config** file
- the **Connection** registry key
- the **Failover** registry key

If you have (as most do) used the default database administration account when installing Fusion PACS, the user ID and password are stored in the clear in the three places listed above. If the customer's Web server were somehow compromised, a hacker could do serious damage with full access to the SQL server databases.

To address this potential security weakness, Fusion PACS includes a utility that performs two functions:

- it creates a new database user with permissions restricted to modifying the two FUSION databases, and changes the web.config file and registry to access the database using this new user account
- it can encrypt the database user ID and password so that it no longer appears in the web.config file or the registry in comprehensible form

Creating a new database user

Note: This step is only necessary if you set up Fusion PACS to use the database administrator account for database access.

This information is stored in several places in the system where someone might find it. You can, however, use a utility included with Fusion PACS to set up a new database user and change the Fusion PACS configuration to access the database with the new user ID. This utility can also be used to change an existing Fusion PACS database user.

This new user ID is restricted to making changes to the two Fusion PACS databases, and cannot change other databases or add users.

Encrypting the database access information

If you choose to encrypt the database access information, the user ID and password in the web.config file and Windows registry keys will be replaced with encrypted information. This prevents unauthorized persons from learning your database access information, but also means that if you forget the user ID and password, you can't learn that information from looking at the registry.

Important: You should record the Fusion PACS database user ID and password in Siebel after encrypting the information, so that other service personnel can maintain the site. Under **no** circumstances should you store the user ID and password on the client server.

Running the utility

The FusionServerUserMaintenance utility is automatically installed with the Fusion PACS software. It should be run on the web server (to update the web.config file) and on each machine that is running the FUSION Services (to update the Windows registry).

Note: If the web server is on a separate server, the FUSION runtime services (including this utility) may not exist on the web server. In this case, follow the procedure described in "Updating a stand-alone web server" on page 41 to update the web.config file.

To run the FusionServerUserMaintenance utility (overall procedure):

1. Start the utility and enter the database administration password.
2. Choose which changes to make (create/update the FUSION database user, encrypt the SQL connection string).
3. Choose which files to update (web.config, Windows registry).

To start the FusionServerUserMaintenance utility:

1. On the server you want to update, open a command window and change directories to the Fusion PACS installation directory (by default, `C:\Program Files\Merge eFilm\Fusion Server`).
2. Type `FusionServerUserMaintenance` and press **Enter**. The utility asks whether you are updating a primary or backup server.
3. Type `P` and press **Enter**. The utility asks for the name or IP address of the primary SQL server.
4. Enter the name (or IP address) and press **Enter**.
5. Enter the database administrator password and press **Enter**. The utility will attempt to connect to the SQL server.

To select your changes:

6. Enter the user ID of the Fusion PACS database user. If you are running the utility for the first time, press **Enter**.
7. Type **Y** and press **Enter** to create a new Fusion PACS database user:
   a) Enter a name for the Fusion PACS database user and press **Enter**.
   b) Enter a password for the Fusion PACS database user and press **Enter**.
8. Type **Y** and press **Enter** to encrypt the SQL connection string.

To update the configuration files:

9. Type **Y** and press **Enter** to update the web.config file:
   a) Type the full path to the Fusion PACS Web directory and press **Enter** (by default, `C:\Inetpub\wwwroot\FusionServer`).

   Note: Only choose this option if you are updating the Web server. If the Web server is not running FUSION Services, you must follow the procedure described in "Updating a stand-alone web server" on page 41.

10. Type **Y** and press **Enter** to update the Connection and Failover keys in the registry.
11. The utility will ask you whether you have a backup server. If you say yes, it will instruct you to run the utility on the backup database server.
12. Press **Enter** to quit.
13. Repeat the process on each machine that is running the FUSION Services.

Updating a stand-alone web server

If the web server is not also running the FUSION Services, the FusionServerUserMaintenance utility will not be available on the web server. To update the web.config file, complete the following procedure.

To update the web.config file:

1. Copy the **web.config** file from the web server to a machine that is running the FUSION Services. Make a note of the directory where you stored the web.config file.
2. On the services machine, follow the procedure described in "Running the utility" on page 40.
3. In **step 9a**, enter the local directory where you stored the copy of the web.config file.
4. Once you have finished running the utility, copy the **web.config** file back to the web server.

Setting up load balancing

Fusion PACS uses the Coyote Point Equalizer™ version 5.2 load-balancing traffic management system.

A load-balanced Fusion PACS system provides a preferred solution to a stand-alone server by allowing servers to share loads and providing redundancy, in case of failure, through the use of multiple servers performing the same functions.

The equalizer manages traffic by simply routing requests to the server that is loaded the least. The equalizer assigns requests intelligently, allowing each server to operate at full capacity, instead of some servers remaining idle while others become overloaded. If the equalizer detects a server failure, traffic is redirected to an available server automatically. When the failure is corrected, the equalizer automatically uses the now operational server.

Many variations of a load-balanced system can be used, depending on deployment requirements. For complete redundancy, at least two Fusion PACS image management servers, two Coyote Point Equalizers, and databases are required. The Fusion PACS Workstation server can be installed on a server within the cluster or outside the cluster depending on system requirements and volume. For smaller, low volume deployments, we recommend the Fusion PACS Workstation server be deployed within the cluster on a shared server to reduce hardware costs.

Note: The redundancy and performance requirements should be discussed with the customer, Professional Services, and a Senior Service Engineer before a final system design is decided on.

The server cluster should be built in its own network segment and the equalizer bridges between the internal and external networks. The internal network is a private subnet only used by the Fusion PACS system and the external network is the existing hospital subnet. Multiple network cards, in the equalizer, work to handle the routing of the packets from one network segment to the second network segment.

The following figure illustrates a typical dual network deployment, where the 1.99.146.85.0 subnet represents the existing hospital network and the 10.0.0.0 network represents the FUSION System subnet. An additional benefit of this deployment is that the Coyote Point Equalizer acts as a firewall, securing the Fusion PACS network.

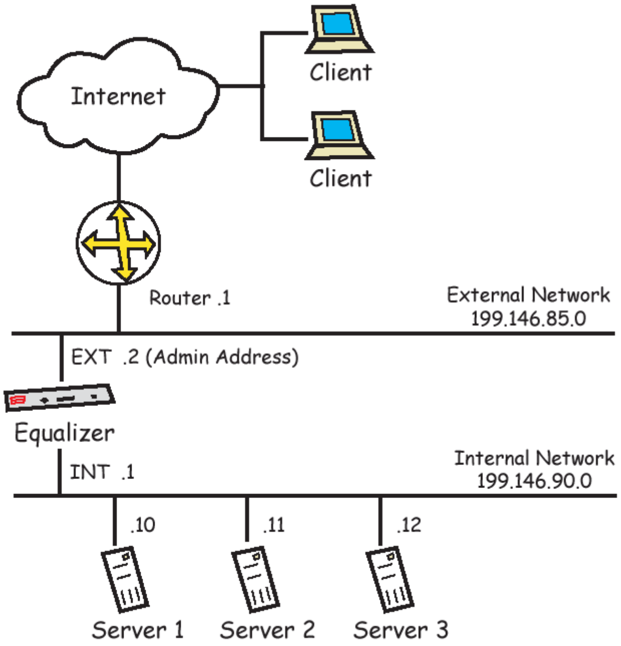

To set up load balancing (overall procedure):

1. Set up the clusters (see “Setting up clusters” on page 43).
2. Build a database for the cluster (see “Building the database for the cluster” on page 43).
3. Build the cluster servers (see “ Building the cluster servers” on page 47).
4. Configure the cluster servers (see “Configuring servers for the clusters” on page 49).
5. Configure the equalizer (see “Configuring the Coyote Point Equalizer” on page 50).
6. Configure the clusters (see “Configuring the clusters” on page 51).
7. Set up a backup equalizer (see “Configuring a backup Coyote Point Equalizer” on page 54).

Setting up clusters

This procedure provides instructions for setting up a cluster.

To set up the cluster:

1. Define an internal private network subnet with the network administrator.

   Tip: This subnet is usually defined as **192.168.X.X**.

2. Reserve an internal IP address to each server within the cluster. This includes any NAS and/or database server.
3. Reserve an external and internal IP address to each Coyote Point Equalizer.
4. Reserve a fail-over gateway from the internal cluster subnet, if dual Coyote Point Equalizers are used.
5. Reserve an external virtual cluster IP address. It will function as the public IP address all modalities send to and clients query to.
6. Install Windows 2003 Server or Advance server on each system in the cluster.
   - The hostname/IP address should be a fixed IP address in a private network subnet.
   - The default gateway must be the internal IP address of the Coyote Point Equalizer, or the fail-over gateway of the Coyote Point Equalizer (in a redundant pair failover Coyote Point Equalizer setup).

   Note: This IP must be on the same subnet as the servers.

   - The DNS must be defined by the Network Administrator.
7. Set up a Windows Workgroup that will include all of the cluster servers.

   Tip: This Windows Workgroup is usually named **MERGE_FUSION**.

   Note: Alternatively, the site may create a separate domain for the Fusion PACS.

8. Define each server name as a unique recognizable name.
9. Define the Administrator username and password as the same for each server to allow services access across servers.

Building the database for the cluster

The database can be installed on a dedicated server or on a server that will also run Fusion PACS. Both types of installation are described in this section. We recommend that the database be run on a dedicated server for clients that run under heavy load or large volume conditions.

Installing the database on a dedicated server

To install the database on a dedicated server (overall procedure):

1. Prepare the server and start the installation (see “Step 1: Starting the installation” on page 44).
2. Install the software for a dedicated installation (see “Step 2: Installing the software” on page 44).

3. Complete the installation (see “Step 3: Completing the installation” on page 45).

Step 1: Starting the installation

To prepare the server and start the installation:

1. Install the following software on the server:
   - Internet Explorer 5.5 or later*
   - IIS 5.0 on each Web server*
   - MSSQL 2000 SP3 or SP3a*
   - Win2000 SP4*
   - Internet Explorer WebControls*
   - .NET Framework 1.1*

   * required to install Fusion PACS 1.5
2. Use Windows Update to verify that you have all the latest service packs installed.
3. Insert the Fusion PACS CD into the CD-ROM drive.

   Note: If you do not have Autorun enabled, explore the CD-ROM and run **launch.exe** to start the Fusion PACS installation program.
4. Click **Installs**, then select **Fusion PACS: Version 1.5**.
5. Click **Next** on the *Welcome* screen and then click **Yes** to accept the license agreement.
6. Click **Next** on the *Information* screen. The *Customer Information* screen appears.
7. Enter the following information, then click **Yes**.
   - Username: **Merge eFilm**
   - Company Name: **Merge eFilm**
   - Select **Install for Anyone who uses this computer**
8. Click **Next**.

Step 2: Installing the software

To build the database on a dedicated server:

1. Select **Fusion PACS Database** and make sure **Fusion PACS Services** and **Fusion PACS Web Site (IIS)** are clear. Click **Next**.
2. Review the information in the *Information* screen, and click **OK**.
3. Click **Next** on the *Available SQL Servers* screen, if the **[local]** database is shown.
4. Add the following information to the *Database* screen, then click **Next**:
   - SQL System Admin: **sa**
   - Password: the sa password assigned when installing the MSSQL Server
   - Computer Name: MSSQL/Windows Computer Name
5. Select **Local Directory Assumed** on the *Database Location* screen, if the default directory was used when installing MSSQL.
6. Click **Next** to accept the default installation destination folder for SQL files.

Step 3: Completing the installation

To complete the installation:

1. Click **Next** to accept the default destination folder for Fusion PACS files.

   Note: We recommend that Fusion PACS files be installed on a volume that has more than 500 MB of free space.

2. Enter the local database server IP address and the FUSION System AE Title, then click **Next**.
3. Click **Next** to accept the default SQL Server IP address and SQL Server name.
4. Click **Next** to accept the setup information.
5. Select the appropriate database installation option:
   - **Create the Necessary FUSION Databases**: will destroy any existing database tables and initialize a new database.
   - **Update the Database**: will add/make any changes and is used when upgrading from an earlier version. This option will not destroy the database tables.
   - **Skip Creating the Database**: will only update stored procedures and is used when reinstalling the Fusion PACS. This option will not affect the database tables.
6. Click **Yes** to continue, then **Finish** to complete the installation.

Installing the database and services on the same machine

To install the database with FUSION services:

1. Prepare the server and start the installation (see "Step 1: Starting the installation" on page 45).
2. Install the software and services (see "Step 2: Installing the software" on page 46).
3. Install the FUSION Services (see "Step 3: Installing the services" on page 47).
4. Conclude the installation (see "Step 4: Finishing the installation" on page 47).

Note: This machine should only be able to be connected to by the Fusion PACS. Proper maintenance and backup routines should be run on this server.

Step 1: Starting the installation

To prepare the server and start the installation:

1. Install the following software on the server:
   - Internet Explorer 5.5 or later*
   - IIS 5.0 on each Web server*
   - MSSQL 2000 SP3 or SP3a*
   - Win2000 SP4* or Windows 2003 Server*
   - Internet Explorer WebControls*
   - .NET Framework 1.1* or ASP.NET* (if using Windows 2003 Server)

   * required to install Fusion PACS 1.5
2. Use Windows Update to verify that you have all the latest service packs installed.

3. Insert the Fusion PACS CD into the CD-ROM drive.

Note: If you do not have Autorun enabled, explore the CD-ROM and run **launch.exe** to start the Fusion PACS installation program.

4. Click **Installs**, then select **Fusion PACS: Version 1.5**.
5. Click **Next** on the *Welcome* screen and then click **Yes** to accept the license agreement.
6. Click **Next** on the *Information* screen. The C*ustomer Information* screen appears.
7. Enter the following information, then click **Yes**.
   - Username: **Merge eFilm**
   - Company Name: **Merge eFilm**
   - Select **Install for Anyone who uses this computer**
8. Click **Next**.

Step 2: Installing the software

To install the database and Fusion PACS services on the same machine:

1. Select **Fusion PACS Services**, **Fusion PACS Web Site (IIS)**, and **Fusion PACS Database**.
2. Click **Next** on the *Fusion PACS–Features* screen to accept a complete Fusion PACS application installation.
3. Click **Next** on the *Available SQL Servers* screen if the **[local]** database is shown.
4. Add the following information to the *Database* screen, then click **Next**.
   - SQL System Admin: **sa**
   - Password: the sa password assigned when installing the MSSQL Server
   - Computer Name: MSSQL/Windows Computer Name
5. Select **Local Directory Assumed** on the Database Location screen, if the default directory was used when installing MSSQL.
6. Click **Next** to accept the default installation destination folder for SQL files.
7. Enter the location where the file system data files will reside on the local computer or on a remote system; for example, **D:\DICOM** or **\\10.1.1.1\DICOM**.

Note: If the file system is installed remotely, file system folders must be shared to allow remote host access.

8. Create a directory for Incoming, Reject, and Storage in the DICOM directory where the file system will be located.
9. Click **Next** to accept the default destination folder for Fusion PACS files.

Note: We recommend that Fusion PACS files be installed on a volume that has more than 500 MB of free space.

10. Enter the local database server IP address and the FUSION System AE Title, then click **Next**.
11. Click **Next** to accept the default SQL Server IP address and SQL Server name.
12. Click **Next** to accept the setup information.

Step 3: Installing the services

To install the services:

1. Select **Install and Set for Activation** on the *FUSION IIS Web Services* screen and click **Next**.
2. Select **Install and Set for Activation** on the *Fusion PACS (Windows NT) Core Services* screen and click **Next**.
3. Select **Specific User – Advanced Setup (Custom)** on the *System User Setup for NT Services* screen and click **Next**.
4. Click **Yes** to continue when informed the specific user must exist and be part of the Administrators group.
5. Enter the Administrator user name and password used to install the operating system.

Note: The same Administrator username and password must be used on all systems running in the FUSION System cluster.

Step 4: Finishing the installation

To conclude the installation:

1. Select the appropriate database installation option:
   - **Create the Necessary FUSION Databases**: will destroy any existing database tables and initialize a new database.
   - **Update the Database**: will add/make any changes and is used when upgrading from an earlier version. This option will not destroy the database tables.
   - **Skip Creating the Database**: will only update stored procedures and is used when reinstalling the Fusion PACS. This option will not affect the database tables.
2. Select **Do Not Activate the Secure IC Connection**, unless users have certificates and want to use the Web IC outside of the firewall. Click **Next** to continue.
3. Click **Finish** to complete the installation.
4. Open Windows Explorer and navigate to **C:\\inetpub\wwwroot\FUSIONServer**. Open the **web.config** file and change the following IP values to the public cluster IP of the Coyote Point Equalizer.
   - `<add key= “cfgMetaDataServerIP” value= “xxx.xxx.xxx.xxx”/>`
   - `<add key= “cfgPixelDataServerIP” value= “xxx.xxx.xxx.xxx” />`

   Note: You will not be able to view images behind the Coyote Point Equalizer following this step.
5. Save and close **web.config**.

Building the cluster servers

This procedure provides instructions for building the cluster servers.

To build the cluster servers:

1. Install the following software (required for Fusion PACS 1.5):

- Internet Explorer 6.0 SP1
- IIS 5.0 on each Web server
- Win2000 SP4 or Windows 2003 Server
- Internet Explorer WebControls
- .NET Framework 1.1 or ASP.NET (if using Windows 2003 Server)

Note: Use Windows Update to verify you have all the latest service packs installed.

2. Insert the FUSION CD into the CD-ROM drive. If you do not have Autorun enabled, explore the CD-ROM and run **launch.exe**. This will start the Fusion PACS installation program.
3. Click **Installs**, then click **Fusion PACS: Version 1.5**.
4. Click **Next** on the *Welcome* screen, then click **Yes** to accept the license agreement.
5. Click **Next** on the *Information* screen.
6. Enter the following information through the *Customer Information* screen, then click **Yes**.
   - Username: **Merge Healthcare**
   - Company Name: **Merge Healthcare**
   - Enable **Install for Anyone who uses this computer**
7. Enable **Fusion PACS Services**, and **Fusion PACS Web Site (IIS)**.
8. Click **Next** on the *Fusion PACS–Features* screen to accept Fusion PACS services and Fusion PACS website installation.
9. Click **OK** to acknowledge that SQL Server is *not* installed locally on the information screen.
10. Add the following information to the *Database* screen, then click **Next**.
    - SQL System Admin: **sa**
    - Password: the sa password assigned when installing the remote MSSQL Server
    - Computer Name: Remote MSSQL/Windows Computer Name
11. Click **OK** to acknowledge SQL Server is *not* installed locally on the information screen.
12. Enter the location where the file system data files will reside on local computer or on a remote system; for example, D:\DICOM or \\10.1.1.1\DICOM.

    Note: If the file system is installed remotely, file system folders must be shared to allow remote host access.

13. Click **Next** to accept the default installation destination folder for Fusion PACS files.

    Note: We recommend Fusion PACS files are installed on a volume that has more than 500 MB of free space.

14. Enter the Fusion PACS IP address and the Fusion PACS AE Title.

    Note: The AE title must be the same for all Fusion PACSs in the cluster.

15. Enter the remote SQL Server IP address and SQL Server name, then click **Next** to continue.

16. Click **Next** to accept the Setup information.
17. Select **Install and Set for Activation** on the *FUSION IIS Web Services* screen and click **Next**.
18. Select **Install and Set for Activation** on the *Fusion PACS (Windows NT) Core Services* screen and click **Next**.
19. Select **Specific User – Advanced Setup (Custom)** the *System User Setup for NT Services* screen and click **Next**.
20. Click **Yes** to continue when informed the specific user must exist and be part of the Administrators group.
21. Enter the Administrator user name and password used to install the operating system.

Note: The same Administrator username and password must be used on all systems running in the FUSION System cluster.

22. Select **Do Not Activate the Secure IC Connection**, unless users have certificates and wish to use the web IC outside of the firewall. Click **Next** to continue.
23. Click **Finish** to complete the installation, and restart the computer.
24. Open Windows Explorer and navigate to **C:\\inetpub\wwwroot\FUSIONServer**. Open the **web.config** file and change the following IP values to the public cluster IP of the Coyote Point Equalizer.
    - `<add key= “cfgMetaDataServerIP” value=”xxx.xxx.xxx.xxx”/>`
    - `<add key= “cfgPixelDataServerIP” value= “xxx.xxx.xxx.xxx” />`

Note: You will not be able to view images behind the Coyote Point Equalizer following this step.

25. Save and close **web.config**.

Configuring servers for the clusters

This procedure provides instructions for configuring the cluster servers.

To configure a cluster server:

1. Connect to **http://localhost/fusionserver** through one of the servers.
2. Log in as **efilm** and enter the password **efilm**.
3. Expand **Configure/Systems/File system**.
4. Set up the File System.

Tip: Remember that the location of this file system must be accessible to all servers and all workstations.

a) Click **Add**.

b) Enter a name for the File System; for example, FUSION Volume 1.

c) Use UNC paths to insert values for DICOM\Storage, DICOM\Incoming, and DICOM/Reject paths; for example, //192.168.0.1/DICOM/Storage

d) Complete the rest of the values as negotiated with the customer.

5. Expand **Configure/Systems/Servers**.

Important: Each machine running the services must be entered through the Add Server window.

6. Add each machine that you want running the services, using the following procedure.
   a) Click **Add**.
   b) Assign a name that identifies the server you are describing.
   c) Enter the hostname for the server.
   d) Enter the **AE Title** that was used during installation.

   Important: Use the same AE Title for each server set up to run services.

   e) Verify that the port is set to **104**.
   f) Enter the value for the Install directory; for example, C:\Program Files\Merge eFilm\Fusion PACS\.
   g) Enter the value for the Log directory; for example, C:\Program Files\Merge eFilm\Fusion PACS\Log.
   h) Click the **Services** Tab.
   i) Add only the following services, with the Log Level set to **Info**:
      - Dicom Service
      - Storage Service
      - Image Channel Service
      - Move Service
      - Media Service*
      - Disk Server*

   Note: * Add these services *only* if they are required to run on the system.

   j) Click **OK**.
7. Repeat step 6 until each of the servers has been added.
8. Select the first server entry in the cluster, and click **Edit**.
9. Select **Service** and add the services that are not built for clustering:
   - Edit Service

Configuring the Coyote Point Equalizer

This section provides the procedure for configuring a Coyote Point Equalizer in a single segment network configuration.

To configure the Coyote Point Equalizer:

1. Plug the network cable into the external NIC card interface.
2. Plug the cables in for power.
3. Connect a serial cable from the serial port of the Coyote Point Equalizer to one of the servers.

Note: The settings are: 9600/8bit/no parity/1 stop bit.

4. On the Windows machine: Configuring HyperTerminal:
   - select **vt100 emulation**
   - select **Properties** from the File menu
   - select the **Settings** tab.
   - enable Terminal keys.
   - set the Emulation mode to VT100.
   - select the **Terminal Setup** button.
   - enable Keypad application mode.
   - enable Cursor keypad mode.
   - set the Character set option to ASCII or Special Graphics.
5. Start up the Coyote Point Equalizer and start the HyperTerminal session.
6. Login as **eqadmin** and enter the password that is provided in the *Equalizer Installation and Administration Guide*.
7. Select **3** – **Interfaces** to configure the external interface for the Coyote Point Equalizer

   a) Select **fxp0 – External Ethernet interface**.
      - configure the Hostname, Domain name, Gateway, Name server, IP Address, and Net Mask
      - click **OK**.

   b) Select **fxp1 – Internal Ethernet interface**.
      - configure the IP Address, and Net Mask
      - click **OK**
8. Select **4** at the Equalizer Configuration Menu to set the time zone.

   Note: Set the time zone for your locality and click **OK**.
9. Select **5** at the Equalizer Configuration Menu to set the clock.

   Note: Set the time according to your location and click **OK**.
10. Select **6** at the Equalizer Configuration Menu to set the Password.

    Note: Set the proper Fusion PACS password and click **OK**.
11. Select **8** at the Equalizer Configuration Menu, to confirm the changes.
12. Reboot the equalizer.

Configuring the clusters

To set up the cluster, you need to know the IP addresses and names of the servers that are to be clustered together. You will also need a Cluster IP address, which will be used by the clients to access the system. This is known as the Virtual IP.

Note: Configuration instructions differ between models and software version of the Coyote load balance servers.

The following Fusion PACS ports must be added as clusters:

DICOM	**104** (or whatever port has been configured for DICOM)
Image Channel	**3333**, or **4444** if using SSL
http (use if not secured)	**80**
https (use if secured)	**443**
Remote Desktop	**3389**

To configure a cluster:

1. From one of the servers, use the browser to connect to the Coyote load balance server by typing its external IP address into the explorer address bar. A login screen appears.
2. Enter the username: **touch** and the password that is provided in the *Equalizer Installation and Administration Guide*.
3. On the **Configure** menu, click **Change Passwords**.
4. Change the password to the standard FUSION Cluster password.
5. Click **Add** and select **Add Cluster** from the pull down menu.
6. Enter the following information for the web server cluster:

Cluster Name	The description of the service cluster use. For example, WebViewer (no spaces or special characters can be used).
Cluster Address	The IP address for external connection to this cluster (Virtual IP).
Protocol	Set to Generic TCP.

7. Click **Next**, to continue.

Port	Set to 80, for html traffic.
Sticky Time	Set to 30.
Load balancing Method	Set to Adaptive.
Load balancing response	Set to 5 – Medium.
Enable Inter-Cluster Sticky	Set to Enabled (checked).
Server Agent Port	Set to 0.
ACV Probe String	Make sure the value is blank.
ACV Response String	Make sure the value is blank.

8. Click **Add**.
9. Locate the new entry for this cluster (e.g., Cluster WebViewer TCP) located on the left side of the screen, and select it.
10. Select **Menu** in the configuration information located on the right side of the screen, and select **Add Server** from the drop-down list.

11. Enter the IP address and Hostname of one of the servers that is running the service and click **Add**.
12. Repeat steps 10 and 11 until all the servers in the cluster are added.
13. Add a new cluster by selecting **Add Cluster** from the **Add** drop-down list.
14. Enter the following information for the DICOM cluster:

Cluster Name	The description of the service cluster use. For example, DICOM.
Cluster Address	The IP address for external connection to this cluster (Virtual IP).
Protocol	Set to Generic TCP.

15. Click **Next**, to continue.

Port	Set to 104, for DICOM traffic (or whatever the configured DICOM port is).
Protocol	Set to TCP.
Sticky Time	Set to 0.
Load balancing Method	Set to Adaptive.
Load balancing response	Set to 5 – Medium.
Enable Inter-Cluster Sticky	Set to Disabled (unchecked).
Server Agent Port	Set to 0.
User server Agent	Set to Disabled (unchecked).
ACV Probe String	Make sure the value is blank.
ACV Response String	Make sure the value is blank.

16. Click **Add**.
17. Locate the new entry for this cluster (e.g., Cluster DICOM TCP) located on the left side of the screen, and select it.
18. Select **Menu** in the configuration information located on the right side of the screen, and select **Add Server** from the drop-down list.
19. Enter the IP address and Hostname of one of the servers that is running the service and click **Add**.
20. Repeat steps 18 and 19 until all the servers in the cluster are added.
21. Add a new cluster by selecting **Add Cluster** from the **Add** drop-down list.
22. Enter the following information for the DICOM cluster:

Cluster Name	The description of the service cluster use. For example, ImageChannel.
Cluster Address	The IP address for external connection to this cluster (Virtual IP).
Protocol	Set to Generic TCP.

23. Click **Next**, to continue.

Port	Set to 3333, for HTML traffic (4444 if using SSL)
Protocol	Set to TCP.
Sticky Time	Set to 0.
Load balancing Method	Set to Adaptive.
Load balancing response	Set to 5 – Medium.
Enable Inter-Cluster Sticky	Set to Disabled (unchecked).
Server Agent Port	Set to 0.
User server Agent	Set to Disabled (unchecked).
ACV Probe String	Make sure the value is blank.
ACV Response String	Make sure the value is blank.

24. Click **Add**.
25. Locate the new entry for this cluster (e.g., Cluster Image Channel) located on the left side of the screen, and select it.
26. Select **Menu** in the configuration information located on the right side of the screen, and select **Add Server** from the drop-down list.
27. Enter the IP address and Hostname of the main Fusion PACS and click **Add**.

Configuring a backup Coyote Point Equalizer

This section provides the procedure for configuring a second Coyote Point Equalizer for use as a hot backup.

To configure a backup Coyote Point Equalizer (overall procedure):

1. Set up the console emulator and log in.
2. Configure the interface settings.
3. Set the date and time.
4. Change the password and save your changes.
5. Log in to the administration application.
6. Set the cluster password.
7. Configure the failover settings.

To set up the console emulator and log in:

1. Plug the network cable into the external NIC card interface.
2. Plug the cables in for power.
3. Connect a serial cable from the serial port of the Coyote Point Equalizer to one of the servers.

Note: The settings are: 9600/8bit/no parity/1 stop bit.

4. On the Windows machine, configure the HyperTerminal:

   a) Select vt100 emulation.
   b) On the **File** menu, click **Properties**.
   c) Click the **Settings** tab, and enable Terminal keys.
   d) Change the Emulation mode to **VT100**, and click **Terminal Setup**.
   e) Enable the **Keypad application** and **Cursor keypad** modes.
   f) Change the Character set option to either **ASCII** or **Special Graphics**.

5. Start up the Coyote Point Equalizer and start the HyperTerminal session.
6. Log in as **eqadmin** and enter the password provided in the *Equalizer Installation and Administration Guide*.

To configure the interfaces:

1. On the **Equalizer > Configure** menu, select **3** > **Interfaces** to configure the external Interface for the Coyote Point Equalizer.
2. Select **fxp0 > External Ethernet Interface**.
3. Configure the Hostname, Domain name, Gateway, Name server, IP Address, and Net Mask, and select **OK**.
4. Select **fxp1 – Internal Ethernet interface**.
5. Configure the IP Address and Net Mask, and select **OK**.

To set the date and time options:

1. On the **Equalizer > Configure** menu, select **4** to set the time zone for your locality, and select **OK**.
2. On the **Equalizer > Configure** menu, select **5** to set the time according to your location, and select **OK**.

To change the password and save your configuration:

1. On the **Equalizer > Configure** menu, select **6** to set the Fusion PACS password, and select **OK**.
2. On the **Equalizer > Configure** menu, select **8** to confirm the changes.
3. Reboot the equalizer.

To log in to the Coyote administration application:

1. From one of the servers, use the browser to connect to the Coyote load balance server by typing its external IP address in the address bar. A login screen appears.
2. Enter the username **touch** and the password that is provided in the *Equalizer Installation and Administration Guide*.
3. Click **OK**.

To set the cluster password:

1. On the **Equalizer > Configure** menu, click **Change Passwords**.
2. Change the password to the standard Fusion PACS Cluster password.

Configuring the failover settings

The following procedure should be followed for both the primary and backup equalizers.

To configure the failover settings:

1. On the **Equalizer > Configure** menu, click **Failover settings**.
2. Select **Dual Network** from the **Configuration Type** drop-down menu.
3. Select the appropriate role for the equalizer you are configuring from the **Failover Role** drop-down menu.
4. In the **Sibling External Address** field, enter the external IP address of the other equalizer.
5. In the **Sibling Internal Address** field, enter the internal IP address of the other equalizer (only in dual network configurations).
6. In the **Failover Gateway Address** field, enter an additional IP address.

   Note: The Failover Gateway IP address will be used as the default gateway on the cluster servers.

7. Click **Set (reboot)** on the primary equalizer first to confirm the changes and reboot the equalizer. Rebooting the primary equalizer first ensures that it will assume the primary role.
8. Reboot the backup equalizer.

Note: When both equalizers are up and running, the cluster configuration settings from the primary Coyote should be automatically sent to the backup Coyote.

Configuring the Matrix Interface

3

This chapter provides instructions for configuring the interface to the Fusion Matrix workstation software, which allows Fusion PACS to act as the image archive for Matrix.

Note: This section assumes that you have installed the Fusion Matrix software on the same server as Fusion PACS. For assistance with other configurations of the software (for example, running multiple instances of the image converter, or spreading the components among several machines), please contact Merge Healthcare Professional Services.

This chapter is divided into the following sections:

- Overview: describes the components of the Fusion Matrix interface.
- Setting up the Cache Manager database and directory
- Configuring the interface software: describes how to configure the MFS component
- Configuring the image converter: describes how to configure the MFICM component
- Configuring Matrix: describes how to configure Matrix to use Fusion PACS as its image archive
- Registering the services: describes how to register the MFS and MFICM services
- Setting up key images: describes how to configure Fusion Server to handle key images from Matrix.

Overview

The Fusion Matrix interface software allows Fusion PACS to act as the image archive for the Fusion Matrix workstation software.

The Fusion Matrix interface consists of three components:

- MFS: Matrix Fusion Server
- MFICM: Matrix Fusion Image Conversion Module
- Utilities: Consists of two utilities, Converter and Retriever.

Setting up the Cache Manager database and directory

You must manually create the Cache Manager database and image directory on a server that can be accessed by the MFS and MFICM modules.

To set up the Cache Manager:

1. In Enterprise Manager, manually create the **FusionServerCacheManager** database on the same server as Fusion PACS.
2. In SQL Query Analyzer, run **SanFran.sql** to populate the FusionServerCacheManager database and the InstanceInfo table in the FusionServer database.
3. Manually create a cache directory; for example: **D:\Cache**.

Configuring the interface software

To configure the Matrix Fusion Server component, edit the **MFS.exe.config** file, located in the C:\Program Files\Merge eMed\MFS directory by default.

You **must** edit the following keys:

Key	Description
ConnectionString	The connection information for the FusionServer database: ■ server: the IP address of the machine hosting the FusionServer database ■ uid: a user ID that has read/write access to the FusionServer database ■ pwd: the password for the database account ■ database: The name of the database
ConnectionStringCache	The connection information for the FusionServerCacheManager database: ■ server: the IP address of the machine hosting the FusionServerCacheManager database ■ uid: a user ID that has read/write access to the FusionServerCacheManager database ■ pwd: the password for the database account ■ database: The name of the database
imagesIncomingDir	Folder for incoming DICOM images. Must be expressed as an IP address followed by the path. **Note:** You may need to create this path as a shared folder on the target machine such that the UNC path exists.
imagesStorageDir	Folder where DICOM images are stored. Must be expressed as an IP address followed by the path. **Note:** You may need to create this path as a shared folder on the target machine such that the UNC path exists.
matrixServerAddress	IP address of the Matrix Server.

Key	Description
cacheImagesRootDir	Folder for cache images. Must be expressed as an IP address followed by the path. **Note:** You may need to create this path as a shared folder on the target machine such that the UNC path exists.
cacheManagerAddress	IP address of the machine where the Cache Manager is running.

Configuring the image converter

You can run more than one instance of the image converter, either on the same machine or on other machines. If you are running multiple instances of the image converter, one must be designated the master instance and all others configured as slaves.

Tip: The master instance does not itself need to convert images; thus, you can set one instance exclusively as a controller and leave image conversion to the slave instances.

For each instance of the image converter that you have created, you must manually edit the **MFICM.exe.config** file, located in the C:\Program Files\Merge eMed\MFICM directory by default.

You **must** edit the following keys:

Key	Description
path2tskill	Sets the path to the tskill.exe file. Depending on your operating system, this file is in either the C:\winnt\system32\ or C:\windows\system32\ folder. Locate the file on your system and set the key value appropriately.
sqlConnectionStringData	The connection information for the FusionServer database: ▪ server: the IP address of the machine hosting the FusionServer database ▪ uid: a user ID that has read/write access to the FusionServer database ▪ pwd: the password for the database account ▪ database: The name of the database
sqlConnectionStringCache	The connection information for the FusionServerCacheManager database: ▪ server: the IP address of the machine hosting the FusionServerCacheManager database ▪ uid: a user ID that has read/write access to the FusionServerCacheManager database ▪ pwd: the password for the database account ▪ database: The name of the database

Key	Description
cacheImagesRoot	Location of the cache directory. Must be specified as a UNC path (i.e., no hostnames are allowed). **Note:** You may need to create this path as a shared folder on the target machine such that the UNC path exists.
cacheSoftLimitMB	The size in MB to which the image cache can grow before the image converter starts purging images.
cacheHardLimitMB	This value should be set at least 512 MB higher than the soft limit, because under no circumstances do you want the cache to reach the hard limit.
numHighPriorityConversionThreads	This key should only be modified by Merge Healthcare Professional Services staff.
numLowPriorityConversionThreads	This key should only be modified by Merge Healthcare Professional Services staff.
runCacheManager	Indicates whether this instance of the image converter will act as the Cache Manager. There must be one and only one Cache Manager in each system, regardless of the number of image converters.
runConverter	Indicates whether this instance of the image converter is converting images. You would only set this value to **false** in the event that you wanted the image converter to act as a dedicated Cache Manager.
controllerAddress	The IP address of the machine where this instance of the image converter is running.
controllerPort	The port number on which this instance of the image converter listens for requests.
masterCacheManagerAddress	If runCacheManager is set to **false**, you must specify the IP address of the Cache Manager in this key.

Configuring Matrix

If the Matrix Server is on a different server than the Fusion Matrix software, you must configure two files as shown below.

Editing settings.xml

Edit Matrix's **settings.xml** (usually located in the D:\eMed\FUSION Matrix Server Components v2.2 directory) file as follows:

1. Make sure that the Apollo.Server.Data.Archive.ArchiveProvider is set to the Fusion server's IP:

```
<config entity="Apollo.Server.Data.Archive.ArchiveProvider">
    <connectionString>http://xxx.xxx.xxx.xxx:10000</connectionString>
```

2. Insert the Matrix server's IP in between the `<respondto></respondto>` that is located before the `<port>2323</port>` line
3. Set `<uri>tcp://zzz.zzz.zzz.zzz</uri>` to `127.0.0.1`

Editing hermes.xml

The Hermes.xml file is usually located in the D:\eMed\FUSION Matrix Server Messaging v2.2 directory.

Tip: To verify what IP addresses are listening to what ports you can run `Netstat -a -n`

To edit the Hermes.xml file:

1. Find the section that starts with 'Listen for Archive query results on port 2323'.
2. Set the source type to the Matrix server IP; for example:

```
<InputList routing="broadcast">
   <source type="http">http://zzz.zzz.zzz.zzz:2323/</source>
```

Registering the services

You must register the MFICM and MFS services on the machines where those services are installed.

Registering MFS

To register MFS:

1. Open a command window.
2. Navigate to the **C:\Program Files\Merge eMed\MFS** folder.
3. Enter the following command and press **Enter**: `wscript register.vbs`
4. Enter the following command and press **Enter**: `register_service.bat`
5. Open the Windows Services widget (right-click **My Computer** and select **Manage**, then expand **Services and Applications > Services**).
6. Right-click the **MFS** service and select **Properties**.
7. Ensure that the Startup Type is set to **Automatic**.
8. Start the service.

Registering MFICM

To register MFICM:

1. Open a command window.
2. Navigate to the **C:\Program Files\Merge eMed\MFICM** folder.
3. Enter the following command and press **Enter**: `wscript register.vbs`
4. Enter the following command and press **Enter**: `register_service.bat`
5. Open the Windows Services widget (right-click **My Computer** and select **Manage**, then expand **Services and Applications > Services**).
6. Right-click the **MFICM** service and select **Properties**.
7. Ensure that the Startup Type is set to **Automatic**.
8. Start the service.

Setting up key images

To set up key images:

1. Navigate to **C:\Program Files\Merge eMed\KeyImageTranslator**
2. Open the **KeyImageTranslator.xml** file.
3. Configure the Data Source, User ID and password for your Visualization Services server database.
4. Make a note of the **AETitle** (GSPS_XLATOR) and **ListenPort** (105), as you will need them later
5. In the C:\Program Files\Merge eMed\KeyImageTranslator folder, double-click **register_service.bat**. A popup window will appear.
6. Enter the user and password that you want the service to run under. You must precede the user name with ".\"; for example: `.\Administrator`

   Note: This should be set to the same user as the Fusion DICOM service.

7. Open the Fusion GUI.
8. Create a Device:
   a) Navigate to **Configure/System/Devices**
   b) Click **Add**.
   c) Enter the following:
      - Name: KI
      - AE_Title: GSPS_XLATOR
      - Port: 105
      - Hostname: The host name of the Fusion Server where the KeyImageTranslator is installed
      - IP Address: The IP address of the Fusion Server where the KeyImageTranslator is installed
      - Select the **Device Enabled** check box.
      - Select the **Enable AutoRouting of Non-Image Objects** check box
      - Select the **Enable Query/Retrieve or Move of Non-Image Objects** check box

   d) Click **OK**.
9. Create a Routing Rule:
   a) Navigate to **Configure/Rules/Routing**
   b) Click **Add**.
   c) Locate the new device under Available Devices and move it to Selected Devices.
   d) Select the **Condition** tab.
   e) Click **Add Condition**.
   f) Create the condition **Modality = 'PR'**.
   g) Click **OK**.
   h) Select the **Action** tab.
   i) Select **Move**.
   j) Enter a name in the **Rule Name** field; for example, Move PR.

k) Click **OK**.

10. Using Windows Explorer, navigate to **C:\Inetpub\wwwroot\FusionServer**.
11. Open the **web.config** file.
12. Navigate to the **<!-- Key Images Configuration -->** section
13. Set the **KeyImagesServerEnabled** key to "true".
14. Set the **KeyImagesServerUrl** value to your Visualization Services server; for example, `http://10.1.33.161/VisualizationServices/KeyImages.asmx?wsdl`
15. Save and close the file.

Utilities

4

This section describes how to run the various utilities that are included with Fusion PACS. The following utilities are documented here:

- Audit log export utility (see“Exporting the audit log” on page 65)
- Database access encryption utility (see “Securing database access” on page 66)
- Tape recovery utility (see “Using the Tape Recovery Utility” on page 68)

Exporting the audit log

Left to itself, the audit log will accumulate entries indefinitely, ulimately compromising system performance. A utility is included with Fusion PACS that will export entries from the audit log to a file and then delete the exported entries from the database. This utility can be configured to run automatically at scheduled intervals.

Required files

The audit log export utility is installed in the Fusion PACS **runtime** directory on any machine that is running the FUSION Services. However, the utility must be run on the machine where the SQL Server database is installed. If the database is running on a stand-alone machine without the FUSION Services, you must copy the following files from a FUSION Services machine to the database machine:

- hipaa_export.bat
- hipaa_rename.vbs

Running the utility

On the database server (see “Required files” above), navigate to the FUSION Server runtime directory (or wherever you placed the required files) and enter the following command:

```
hipaa_export.bat username password export_delay
```

...where:

- `username` is the user ID for the database administration account
- `password` is the password for the database administration account
- `export_delay` sets a period, in days, beyond which entries will be exported and removed from the database; for example, if you set the export delay to 7, then all audit log entries older than a week will be removed from the database and exported to disk

Viewing the exported file

Exported audit log entries are copied to a comma-separated value file in the same directory as the utility. The file is named according to the following convention:

```
hipaa_log_YYYYMMDD_HHMMSS.csv
```

...where `YYYYMMDD` is the date and `HHMMSS` is the time the file was created. The exported file can be viewed in any spreadsheet program.

Securing database access

By default, Fusion PACS stores the database user ID and password in three places:

- the **web.config** file
- the **Connection** registry key
- the **Failover** registry key

If you have (as most do) used the default database administration account when installing Fusion PACS, the user ID and password are stored in the clear in the three places listed above. If the customer's Web server were somehow compromised, a hacker could do serious damage with full access to the SQL server databases.

To address this potential security weakness, Fusion PACS includes a utility that performs two functions:

- it creates a new database user with permissions restricted to modifying the two FUSION databases, and changes the web.config file and registry to access the database using this new user account
- it can encrypt the database user ID and password so that it no longer appears in the web.config file or the registry in comprehensible form

Creating a new database user

Note: This step is only necessary if you set up Fusion PACS to use the database administrator account for database access.

This information is stored in several places in the system where someone might find it. You can, however, use a utility included with Fusion PACS to set up a new database user and change the Fusion PACS configuration to access the database with the new user ID. This utility can also be used to change an existing Fusion PACS database user.

This new user ID is restricted to making changes to the two Fusion PACS databases, and cannot change other databases or add users.

Encrypting the database access information

If you choose to encrypt the database access information, the user ID and password in the web.config file and Windows registry keys will be replaced with encrypted information. This prevents unauthorized persons from learning your database access information, but also means that if you forget the user ID and password, you can't learn that information from looking at the registry.

Important: You should record the Fusion PACS database user ID and password in Siebel after encrypting the information, so that other service personnel can maintain the site. Under **no** circumstances should you store the user ID and password on the client server.

Running the FusionServerUserMaintenance utility

The FusionServerUserMaintenance utility is automatically installed with the Fusion PACS software. It should be run on the web server (to update the web.config file) and on each machine that is running the FUSION Services (to update the Windows registry).

Note: If the web server is on a separate server, the FUSION runtime services (including this utility) may not exist on the web server. In this case, follow the procedure described in "Updating a stand-alone web server" on page 68 to update the web.config file.

To run the FusionServerUserMaintenance utility (overall procedure):

1. Start the utility and enter the database administration password.
2. Choose which changes to make (create/update the FUSION database user, encrypt the SQL connection string).
3. Choose which files to update (web.config, Windows registry).

To start the FusionServerUserMaintenance utility:

1. On the server you want to update, open a command window and change directories to the Fusion PACS installation directory (by default, `C:\Program Files\Merge eFilm\Fusion Server`).
2. Type `FusionServerUserMaintenance` and press **Enter**. The utility asks whether you are updating a primary or backup server.
3. Type `P` and press **Enter**. The utility asks for the name or IP address of the primary SQL server.
4. Enter the name (or IP address) and press **Enter**.
5. Enter the database administrator password and press **Enter**. The utility will attempt to connect to the SQL server.

To select your changes:

6. Enter the user ID of the Fusion PACS database user. If you are running the utility for the first time, press **Enter**.
7. Type **Y** and press **Enter** to create a new Fusion PACS database user:
   a) Enter a name for the Fusion PACS database user and press **Enter**.
   b) Enter a password for the Fusion PACS database user and press **Enter**.
8. Type **Y** and press **Enter** to encrypt the SQL connection string.

To update the configuration files:

9. Type **Y** and press **Enter** to update the web.config file:
   a) Type the full path to the Fusion PACS Web directory and press **Enter** (by default, `C:\Inetpub\wwwroot\FusionServer`).

Note: Only choose this option if you are updating the Web server. If the Web server is not running FUSION Services, you must follow the procedure described in "Updating a stand-alone web server" on page 68.

10. Type **Y** and press **Enter** to update the Connection and Failover keys in the registry.
11. The utility will ask you whether you have a backup server. If you say yes, it will instruct you to run the utility on the backup database server.
12. Press **Enter** to quit.
13. Repeat the process on each machine that is running the FUSION Services.

Updating a stand-alone web server

If the web server is not also running the FUSION Services, the FusionServerUserMaintenance utility will not be available on the web server. To update the web.config file, complete the following procedure.

To update the web.config file:

1. Copy the **web.config** file from the web server to a machine that is running the FUSION Services. Make a note of the directory where you stored the web.config file.
2. On the services machine, follow the procedure described in "Running the FusionServerUserMaintenance utility" on page 67.
3. In **step 9a**, enter the local directory where you stored the copy of the web.config file.
4. Once you have finished running the utility, copy the **web.config** file back to the web server.

Using the Tape Recovery Utility

The FUSION Tape Recovery Utility is a command line tool that allows you to restore DICOM series archived through the FUSION Media Service from FUSION version 1.2 through version 1.5. The utility is intended for use in situations where your site has migrated away from Fusion PACS (or away from an earlier version of Fusion PACS created a particular tape) or in the event that a tape has somehow become damaged.

The utility places all recovered data into a directory structure of study and series subdirectories that contains standard DICOM part-10 explicit little endian encoded images. It also creates a log file in the execution directory that contains a list of the restored series.

As the tool is designed to restore the entire tape, you will be unable to pick a specific series to restore. The tool will do its best to deal with faulty tapes: it will begin by checking the beginning and the end of the tape; it will then attempt to find the first available series archive. If it is unable to read it, it will keep skipping through the tape trying to recover any data it can find.

Note: Execution of this tool requires a computer running Windows 2000 or later and exclusive access to the tape drive during the time of its operation.

This section describes how to:

- install the utility (see "Installing the utility" on page 69)

- run the utility (see "Using the utility" on page 69)
- review the log file ("Reviewing the Tape Recovery Log" on page 71)

Installing the utility

Included with the FUSION 1.5 install in the **service** subdirectory is a ZIP file named **TapeRecoveryUtil.zip**. You will need to copy this zip file to the machine attached to the tape drive you will be using and extract the files inside to an empty folder. No configuration is required for the utility.

WARNING! Do **not** extract the utility to the FUSION Server directory.

Using the utility

To recover image files using the Tape Recovery Utility (high-level procedure):

1. Determine the number of the drive you are using (see "Determining the drive number" on page 69).
2. The second step depends on whether you are recovering from a tape drive or a NAS file system:
   - **If you are recovering from a tape drive**, mount the tape from which you want to extract the data (see "Mounting the tape" on page 69).
   - **If you are recovering from a file system**, map a network drive on the machine where you are running the utility to the root directory of the archive file system.
3. Run the utility (see "Running the utility" on page 70).

Determining the drive number

To determine the drive's ID:

1. Right-click the drive in the "Physical Locations" section of the Windows RSM, and select **Properties**.
2. Select the **Device Info** tab. You should see the "Device Name" somewhere on that screen (depending on what version of Windows you are running). It will look something like \\.\Tape0. The number after Tape is the drive ID.

Mounting the tape

This step is only required if you are recovering data from a tape system. If you are recovering data from a NAS file system, you may skip this step.

To mount the tape:

1. If the drive exists on the system on which the FUSION Media Service is running, stop the service during the restoration.

   Tip: Since the tape restoration process can take quite some time, and since the drive can't be used for any other purpose while the utility is running, we recommend that you use a different system for the recovery, providing it has an equivalent and compatible tape drive.

2. If using a library or autoloader, load the source tape.
3. Mount the tape into the drive using the Windows Removable Storage Manager (RSM).

Running the utility

Once started, the utility will run until it has restored every series it can access on the tape, or it encounters a critical error that it can't skip over. This may take anywhere from a few minutes to 12 hours or more, depending on the hardware, state of the tape, and quantity of data. So be certain you don’t need the machine for anything else for that period.

Important: For this reason, we suggest that you do not run the utility on the archive server.

To run the tape recovery utility:

1. From the directory where you extracted the zip file, execute **TapeRecoveryUtility.exe**.
2. The tool will check for an existing log file. If one exists, it will ask you whether you want to overwrite it or not.

   Note: If you choose **No**, execution will cease and you will need to manually move or rename the existing log file.

3. The tool will then display the following prompt,

   ```
   Starting Merge eFilm Fusion PACS Tape Recovery Utility. Please STOP
   FUSION MEDIA SERVICE immediately.


   Please select one of the following modes (enter 0 or 1), and press
   ENTER.


   0 (Tar file recovery from tape)
   1 (Recovery from mounted directory: NetAttached (StorageTEK) mode.)
   ```

   If you are recovering from tape, enter **0** and press **Enter**, then proceed to step 4.

   If you are recovering from a file system, enter **1** and press **Enter**, then proceed to step 9.

4. If you are recovering from tape, the tool displays the following prompt:

   ```
   Insert the tape to recover into the appropriate drive, mount it with
   the Removable Storage Manager, and Press Enter to continue.
   ```

   Assuming you followed the procedure under “Mounting the tape” on page 69, press **Enter**.

5. The tool will then prompt:

   ```
   What is the drive number? (Drives start at 0) :
   ```

   Enter the value you determined in “Determining the drive number” on page 69.

6. The utility will now attempt to access this tape.

   If it is unable to find the drive or access the tape, it will display an error and wait for you to exit the application. If it is unable to access the drive it is likely that the drive ID is incorrect, or the drive is currently in use by another process. If it is unable to access the tape, you can check the RSM's status log to see if the system is currently busy performing some other operation.

7. The utility will now check the beginning and end of the tape.

If the standard Microsoft Tape Format (MTF) header is intact and available, the utility will display this information. You will not need to know any of this information; it's purely for information purposes. If the MTF header is unavailable or corrupt the utility will display:

```
The current tape is either empty, does not contain any TAR files, or
the first file is a not a tar file. Would you like to try to continue?
(y/n) :
```

8. Enter **Y** and press **Enter**, unless you think that you have loaded a blank tape by mistake and do not want to wait for the utility to scan for studies. Proceed to step 10.
9. If you are recovering from a file system, you will be prompted for the directory. Enter the name of the network drive you mapped to the archive file system and press **Enter**.
10. You will now be prompted for the directory to which to restore the series:

    ```
    What is the root directory to copy the study/series to? (Please don't
    use Fusion SD's Incoming directory, as it will create permission
    issues.)
    Directory? (ex: C:\\Recovery\\) :
    ```

    WARNING! Don't restore series directly to any of FUSION's file systems.

11. Confirm your directory choice.
12. Once this is complete, the tool will run through the tape or file system attempting to restore the data.

    For each series, a line will appear saying:

    ```
    Attempting to read archive #N
    ```

    ...where N is the archive number from the start of the tape.

    Note: This is the only display of the tool's status. It will not estimate percent complete. If an error is encountered during extraction of a series, the tool displays `Skipping unknown or corrupt archive` and continues to the next series.

13. At this point, the utility requires no further operator intervention. When execution is completed, you will be prompted to exit the program. The data directory you entered above will contain a separate subdirectory for each of the studies restored. Each study will contain subdirectories for the series restored for that study, and each series directory will contain all the images restored for that series.

Tip: Information about recovered data is recored in the **TapeRecoveryUtil.log** file in the execution directory (see "Reviewing the Tape Recovery Log" on page 71).

Reviewing the Tape Recovery Log

The tape tool logs all restored images to a file named **TapeRecoveryUtil.log**. Each line in the file corresponds to a recovered image. The lines are semicolon delimited, and are formatted as follows:

```
Patient's Name; Patient's ID; Accession Number; Study Date; Modality;
Study UID; Series UID
```

Installing SQL Server Clustering

A

This appendix describes how to install a Microsoft SQL Server Cluster environment for the Fusion PACS Storage and Distribution Manager. This document is divided into the following sections:

- "Before you install" on page 73 describes the prerequisites for setting up SQL Server clustering.
- "Installing the Cluster Service software" on page 81 describes how to install Microsoft Clustering Services, which are required to run a SQL Server Cluster.
- "Installing SQL Server Clustering" on page 90 describes how to install the SQL Server Cluster.

Before you install

This section describes the software and hardware prerequisites for setting up SQL Server Clustering. It also describes the steps you will have to take to set up the networks.

System requirements

There are two types of cluster configurations supported by Microsoft:

- A 2 node cluster includes 2 servers that are each a node in the cluster. This configuration is supported in Windows Advanced Server 2000 and Windows 2000 Datacenter Server.
- A 4 node cluster includes 4 servers that are each a node in the cluster. This configuration is only supported in Windows 2000 Datacenter Server.

Server/Node Hardware: Each node in the tested environment are Dell PowerEdge 2600 2.4 GHz servers. Each node has at least 1 GB of RAM. Each node contains a Dell PERC 4/ DC RAID Controller to connect to the shared storage array. Each node must have 2 Network Interface Cards. 1 NIC on the node will provide Local Network connectivity, the other NIC will be used for a private network, called the heartbeat, in the cluster.

Storage Array: The storage array used for testing is a Dell PowerVault 220S. 5 drives each with 33 GB of storage exists on the array.

Operating System: Microsoft Windows Advanced Server 2000 with all current service packs (2-node cluster only) or Windows 2000 Datacenter Server (2- or 4-node clusters)

System prerequisites

Before attempting to install a Microsoft SQL Server Cluster make sure all of the following requirements have been met:

- Microsoft Windows 2000 Advanced Server, Windows 2000 Datacenter Server, Windows 2003 Enterprise Edition or Windows 2003 Datacenter Edition installed on all computers in the cluster.
- You are using a name resolution method such as DNS, WINS, HOSTS, etc.
- All hardware meets the Windows 2000 Advanced Server, Windows 2000 Datacenter Server, Windows 2003 Enterprise Edition or Windows 2003 Datacenter Edition product compatibility list. This can be located here: http://www.microsoft.com/whdc/hcl/default.mspx
- The cluster hardware must be on the Cluster Service Hardware Compatibility List (HCL). This can be found here: http://www.microsoft.com/whdc/hcl/default.mspx
- The network has:
  - A unique NetBIOS name for the cluster.
  - Four unique static IP Addresses for the network. Two will be used on the nodes for the public network, one will be used for the Windows Cluster, and the last will be used by the SQL Server Virtual Server.
- A Domain Administrator user account that will run all of the Windows Cluster and SQL Server Cluster services.

Setting up the networks

1. Power down all nodes and the shared storage device.
2. Make sure the nodes in the cluster can not access the storage device.
3. Verify that all physical network connections are correct: The public network adapters must be connected to the public network and the private network adapters are connected to each node only

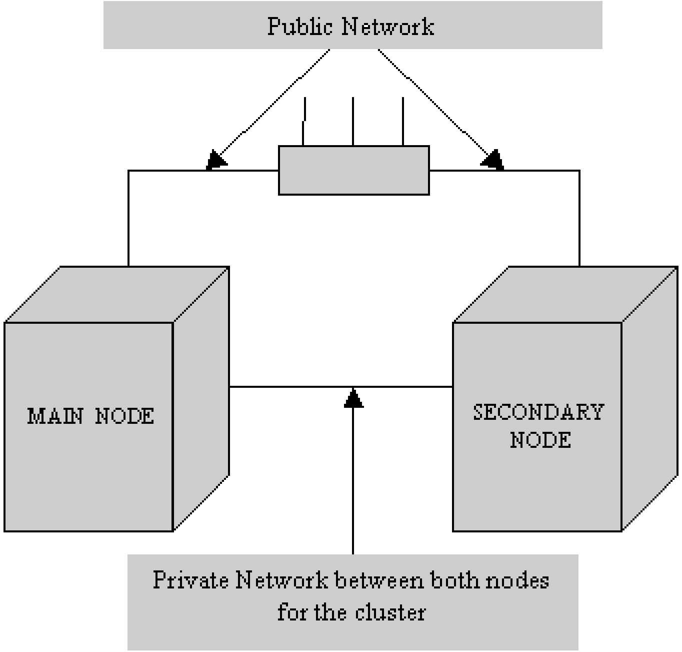

4. When you have verified that all network connections are correct power on both nodes.

Configuring the Private Network Adapter for the Cluster

This section describes the steps in configuring the Private Network Adapter that will be used for the internal cluster communication.

To configure the private network adapter:

1. Select **Start > Settings > Control Panel** and double-click **Network Connections**.
2. Select **Advanced> Advanced Settings...**
3. In the **Connections** box, make sure that your bindings are in the following order, and then click **OK**:
   a) External Public Network
   b) Internal Private Network (heartbeat)
   c) Remote Access Connections
4. Right-click the network connection for your heartbeat adapter and click **Properties**.
5. On the *General* tab, make sure that only the **Internet Protocol (TCP/IP)** check box is selected. Clear all other check boxes.
6. Make sure that both of the private network adapters on each node are set to the same speed.
7. Click **Internet Protocol (TCP/IP)**, and then click **Properties**.

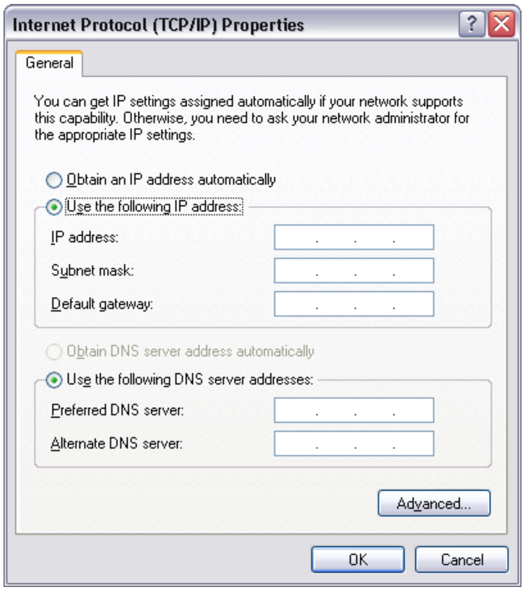

Internet Protocol (TCP/IP) Properties

General

You can get IP settings assigned automatically if your network supports this capability. Otherwise, you need to ask your network administrator for the appropriate IP settings.

○ Obtain an IP address automatically
● Use the following IP address:
IP address: . . .
Subnet mask: . . .
Default gateway: . . .

○ Obtain DNS server address automatically
● Use the following DNS server addresses:
Preferred DNS server: . . .
Alternate DNS server: . . .

8. On the *General* tab, verify that you have selected a static IP address that is not on the same subnet or network as another one of the public network adapters.

   An example of good IP addresses to use for the private adapters is 10.10.10.10 on node 1 and 10.10.10.11 on node 2 with a subnet mask of 255.0.0.0. If your public network uses the 10.x.x.x network and 255.0.0.0 subnet mask please use an alternate private network IP and subnet
9. Make sure that there is no value set in the **Default Gateway** field.
10. Verify that there are no values defined in the **Use the following DNS server addresses** field.

11. Click **Advanced...** and select the **DNS** tab.

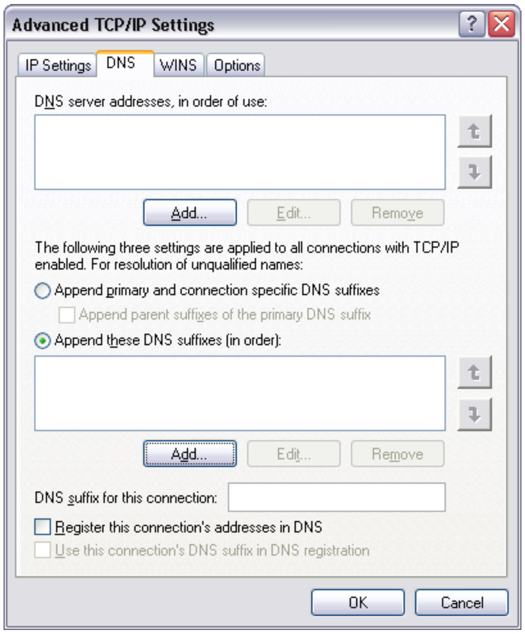

DNS server addresses, in order of use:

The following three settings are applied to all connections with TCP/IP enabled. For resolution of unqualified names:

○ Append primary and connection specific DNS suffixes

□ Append parent suffixes of the primary DNS suffix

● Append these DNS suffixes (in order):

DNS suffix for this connection:

□ Register this connection's addresses in DNS

□ Use this connection's DNS suffix in DNS registration

12. On the DNS tab, verify that there are no values defined. Make sure that the **Register this connection's addresses in DNS** and **Use this connection's DNS suffix in DNS registration** check boxes are cleared

13. On the *WINS* tab, verify that there are no values defined. Select the **Disable NetBIOS over TCP/IP** radio button.

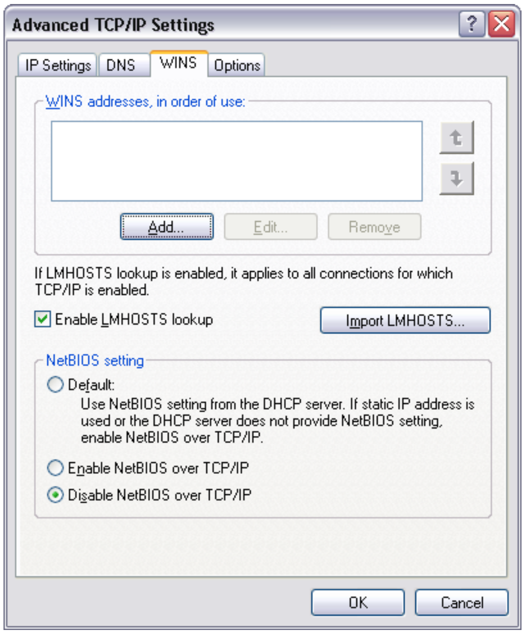

WINS addresses, in order of use:

If LMHOSTS lookup is enabled, it applies to all connections for which TCP/IP is enabled.

☑ Enable LMHOSTS lookup

NetBIOS setting

○ Default:
Use NetBIOS setting from the DHCP server. If static IP address is used or the DHCP server does not provide NetBIOS setting, enable NetBIOS over TCP/IP.

○ Enable NetBIOS over TCP/IP

● Disable NetBIOS over TCP/IP

14. When you close the dialog box, you may receive the following prompt:

```
This connection has an empty primary WINS address. Do you want to
continue?
```

Click **Yes**.

15. You should rename this connection for quick reference. An example of this is to name the connection "Heartbeat".
16. Repeat for all other nodes in the cluster.

Configuring the Public Network Adapter

This section describes the steps required to configure the Public Network Adapter.

To configure the public network adapter:

1. Select **Start > Settings > Control Panel** and double-click **Network Connections**.
2. Right-click the Public Network Adapter and select **Rename**.
3. Give the adapter to a meaningful name for future quick reference. A suggestion is "Public Network"
4. Right-click the Public Network Adapter and select **Properties**.
5. Highlight **Internet Protocol (TCP/IP)** and click **Properties**.
6. Enter the Static IP Address to communicate with the Public Network.
7. Click **OK** twice.
8. Do this for all nodes in the cluster.
9. Restart all nodes when you have configured all the network adapters.

Verifying network connectivity and name resolution

The next steps will help you verify that your Public and Private network connections are working properly.

To verify connectivity:

1. For each node in the cluster, obtain the machine's public and private IP addresses:
   a) Select **Start > Run** to open a command window
   b) Type `cmd` and press **Enter**.
   c) Type `IPConfig /all` and press **Enter**. This will give you all settings for your network adapters.
2. Using a computer on the public network, ping each Public Network Adapter in the cluster using the static IP address for each machine. You should also be able to use the machine names.

   If any of these ping tests fail, verify all of your public network settings and try again.
3. Using a computer on the public network, ping each of the cluster nodes using the Private Cluster Network IP addresses. You should receive timeouts, since these addresses should not be visible from the public network.

   If you are able to communicate with the private cluster network from another computer on the network, verify all of your IP, Subnets, and Default Gateway settings for the Private Network Adapter on both nodes in the cluster.
4. From each node in the cluster, ping the other nodes using the Private Cluster Network (Heartbeat) IP addresses. You should be able to communicateTest the connectivity by running a series of ping tests from both nodes in the cluster and from another computer on the Public Network.

Creating a Cluster user account

This user account will be used for running the Cluster Service. Create a Domain user on the network where the cluster will run. This user should be a Domain Administrator with a password that never expires.

Installing the storage device

To install the storage device:

1. Power down all of the nodes in the cluster and the storage device.
2. Designate one of the nodes to be the "Main" node in the cluster. From this step forward, all references to the "Main" node refer to this node. The all other nodes will be named "Secondary".
3. Attach the SCSI cable from the storage device to the Main node.

WARNING! Be very careful attaching the SCSI cable to the storage device and the Main node, as it is fragile.

4. Power on the storage device. Wait until all lights on the front of the device stop flashing. If any of the lights are not green, then please refer to the manufacturer's user guide for help.
5. Power on the Main node.

Note: In order to successfully configure the storage device, you must **only power on the Main node**. All configurations to the Storage Device and the initial install of the Clustering Service software must be completed with only the Main node attached to the Storage Device and turned on.

6. You might be prompted to install the required drivers for the SCSI card when the Main node starts Windows. At this time, provide the manufacturer's installation drivers for this device.
7. If you have installed a Dell Storage Device, then you should install the Dell OpenManage Software Applications. This package includes a Storage Array manager.

Setting up the Quorum disk

The Quorum disk is used by the cluster to store cluster configuration and database checkpoints. Log files will also be stored on the Quorum disk. The following recommendations should be followed:

- Create a partition of 500 MB on the Storage Device.
- You should dedicate a separate disk for the Quorum and also put this disk on a RAID volume since the cluster, itself, is dependant on this Quorum Disk.

Configuring shared disks

To configure the shared disks:

1. Right-click **My Computer** and select **Manage**. Select Storage.
2. Double click **Disk Management**.

Note: If you are using Dell OpenManage Software then follow the next steps. If you are not, then please refer to your manufacturer’s user manual on how to configure shared disks and setting up RAID 5.

3. Create a Virtual Disk on the RAID Controller:
    a) Click the **Create Virtual Disk** icon.
    b) A window should appear with all of the available RAID controllers on your Main node. Select the proper item and click **OK**.
    c) On the Next Screen you will name your Virtual Disk and set up other properties for it.
    d) Select the **Type** drop-down list and select **RAID 5**.
    e) Select the size you desire.
    f) Select the **Stripe Size**.
    g) For the **Write Policy**, make sure you select **Write-Through**.

    WARNING! Failure to do this will result in an improper Cluster setup.

    h) Select the available disks for the Virtual Disks and click **Finish**.
4. After doing this you will need to Format this virtual disk as **NTFS**.
5. Assign drive letters to each formatted disk on the shared storage device. We recommend assigning the QUORUM disk as drive letter Q. The RAID 5 partition should be assigned a meaningful letter, possibly D for data. If this is not available, then select another meaningful letter.

Note: If you are having problems configuring the shared storage device, then please refer to your manufacturer's user manual on how to do this.

Verifying disk access and functionality

Before continuing, you should verify that you can access the QUORUM disk and the RAID 5 partition from the MAIN Node.

To verify disk access:

1. Open Notepad and create a simple file.
2. Save the file to each shared disk.
3. Try opening the saved file from the shared disks to ensure you can read and write to the disks.

If you are unable to read/write to the shared disks, then you must correct the problem before continuing. Failure to do so will result in an improper Cluster install.

Setting up the Secondary Node

Note: At this time you should shut down the Main node in the cluster. When this node is powered down you can power on the Secondary node.

You will need to ensure that the Secondary nodes RAID controller has a different SCSI ID than the Main node. Failure to do so will create a conflict between the Main and Secondary Nodes when they are both powered on together.

You can assign a SCSI ID at boot time by entering the RAID management utility and selecting the RAID controller. At this time you can assign a different SCSI ID to this controller. By default the SCSI ID is 7; we recommend that you make the Main node's SCSI ID 13 and the Secondary node's SCSI ID 14. This will eliminate the possibility of another default SCSI ID causing a conflict.

If you are using a DELL PERC RAID Controller you must make sure of two things:

- The Write Policy on the controller is Write-Thru
- The Cluster Mode is enabled.

You can view these settings at boot time by accessing the hardware utility program. You need to make sure both nodes have these settings.

After the second node is powered on, you may have to install the RAID Controller drivers. Please use your manufacturer's installation disks to do this.

When the second node is running, repeat the section "Verifying disk access and functionality" on page 79 to ensure that the Secondary node can access the shared storage device.

Note: Once you have completed the above procedures, power down all nodes but leave the shared storage device on.

SCSIdevice issues

If you are having difficulty with your SCSI devices read the following sections to try and help solve the problem.

The SCSI bus listed in the hardware requirements must be configured before installing Cluster services. This includes:

- Configuring the SCSI devices.
- Configuring the SCSI controllers and hard disks to share the SCSI bus.
- Properly terminating the bus. The shared SCSI bus must have a terminator at each end of the bus. It is possible to have multiple shared SCSI buses between the nodes of a cluster.

In addition to the information on the following pages, refer to the documentation from the manufacturer of the SCSI device or the SCSI specifications, which can be ordered from the American National Standards Institute (ANSI). The ANSI web site contains a catalog that can be searched for the SCSI specifications.

Configuring the SCSI Devices

Each device on the shared SCSI bus must have a unique SCSI ID. Since most SCSI controllers default to SCSI ID 7, you will have to to **change the SCSI ID** on one controller to a different SCSI ID, such as SCSI ID 6. If there is more than one disk that will be on the shared SCSI bus, each disk must also have a unique SCSI ID.

Some SCSI controllers reset the SCSI bus when they initialize at boot time. If this occurs, the bus reset can interrupt any data transfers between the other node and disks on the shared SCSI bus. Therefore, **SCSI bus resets should be disabled** if possible.

Terminating the Shared SCSI Bus

There are a couple of methods you can use to terminate the SCSI bus:

- You can connect a Y cable to a device if the device is at the end of the SCSI bus. A terminator can then be attached to one branch of the Y cable to terminate the SCSI bus. This method of termination requires either disabling or removing any internal terminators the device may have.
- Trilink connectors can be connected to certain devices. If a device is at the end of the bus, a trilink connector can be used to terminate the bus. This method of termination requires either disabling or removing any internal terminators the device may have.

Y cables and trilink connectors are the recommended termination methods, because they provide termination even when one node is not online.

Note: Any devices that are not at the end of the shared bus must have their internal termination disabled.

Installing the Cluster Service software

This section describes how to install the cluster service software:

- if you are using one of the Windows 2000 operating systems, see "Installing the cluster service software on Windows 2000" on page 81.
- if you are using one of the Windows 2003 operating systems, see "Installing the cluster service software on Windows 2003" on page 83.

Installing the cluster service software on Windows 2000

Note: At this time you should power on the Main node in the cluster, but leave all other nodes off.

When the Main node is powered up and you are logged into Windows with Local Administrator rights, follow these steps to install the cluster software:

1. Select **Start > Settings > Control Panel > Add/Remove Programs > Add/Remove Windows Components**.
2. Select **Cluster Service** and click **Next**. You will be prompted for the Windows Operating System installation disk.

3. Enter the appropriate path pointing to the version of Windows you are using and the **i386** directory. Click **Next**.

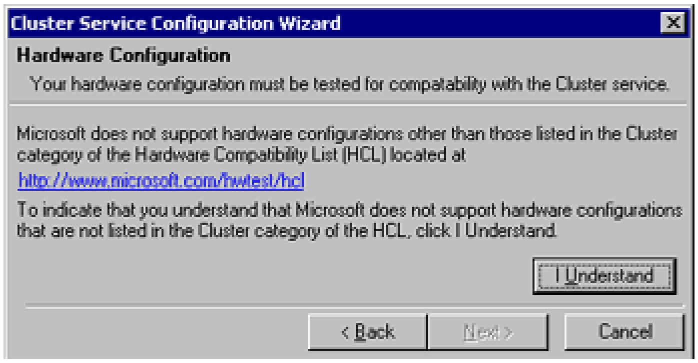

Cluster Service Configuration Wizard

Hardware Configuration

Your hardware configuration must be tested for compatability with the Cluster service.

Microsoft does not support hardware configurations other than those listed in the Cluster category of the Hardware Compatibility List (HCL) located at http://www.microsoft.com/hwtest/hcl

To indicate that you understand that Microsoft does not support hardware configurations that are not listed in the Cluster category of the HCL, click I Understand.

4. The next screen is a Microsoft disclaimer that the equipment you are using in the cluster is on the Hardware Compatibility List. Click **I Understand**, then click **Next**.

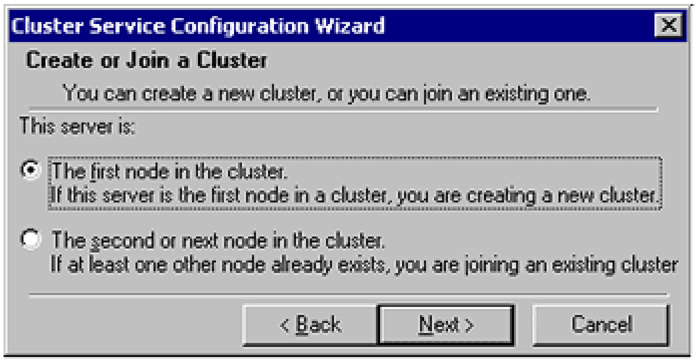

Cluster Service Configuration Wizard

Create or Join a Cluster

You can create a new cluster, or you can join an existing one.

This server is:

● The first node in the cluster.
If this server is the first node in a cluster, you are creating a new cluster.

○ The second or next node in the cluster.
If at least one other node already exists, you are joining an existing cluster

5. Since this is the first node in the cluster, select the **The first node in the cluster** radio button and click **Next**.
6. At the next screen enter a **Name** for the cluster (up to 15 characters) and click **Next**.
7. Type in the Domain Administrator User Account you created in the section "Creating a Cluster User Account" and click **Next**.

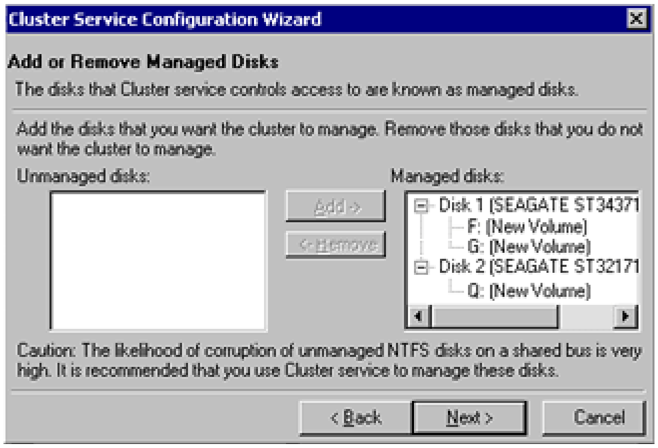

Cluster Service Configuration Wizard

Add or Remove Managed Disks

The disks that Cluster service controls access to are known as managed disks.

Add the disks that you want the cluster to manage. Remove those disks that you do not want the cluster to manage.

Unmanaged disks:

Managed disks:

Disk 1 (SEAGATE ST34371
F: (New Volume)
G: (New Volume)
Disk 2 (SEAGATE ST32171
Q: (New Volume)

Caution: The likelihood of corruption of unmanaged NTFS disks on a shared bus is very high. It is recommended that you use Cluster service to manage these disks.

8. This screen tells the Cluster Service which disks it will be using in the cluster. If there are any disks listed that will not be used in the cluster, make sure you remove them from the managed disks list. Click **Next**.

9. Click **Next** in the *Configuring Cluster Networks* dialog.
10. The next box will show a network connection for the cluster. If the Network adapter listed is the Public Adapter, select the **Enable this network for cluster use** check box.
11. Select the **Client Access Only (public network)** option and click **Next**.

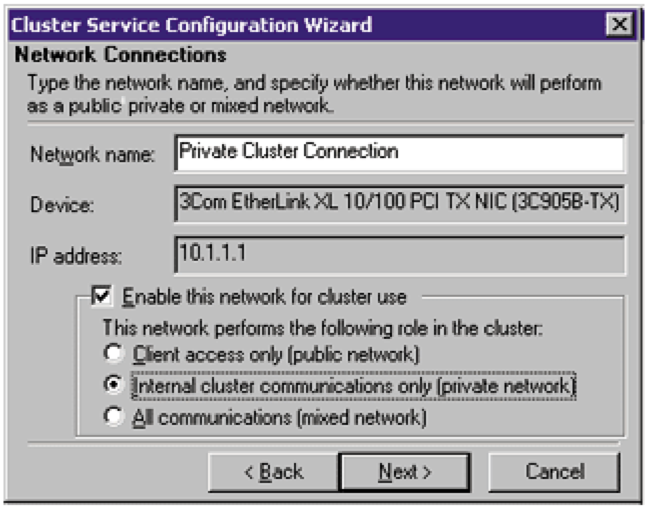

Cluster Service Configuration Wizard

Network Connections

Type the network name, and specify whether this network will perform as a public private or mixed network.

Network name: Private Cluster Connection

Device: 3Com EtherLink XL 10/100 PCI TX NIC (3C905B-TX)

IP address: 10.1.1.1

☑ Enable this network for cluster use

This network performs the following role in the cluster:

○ Client access only (public network)

● Internal cluster communications only (private network)

○ All communications (mixed network)

12. You should receive another dialog for the private network adapter. Select the **Internal Cluster communications only (private network)** option and click **Next**.
13. The next dialog gives the opportunity to order the Network adapters. Make sure the **Private Cluster Connection** is at the top. Click **Next**.

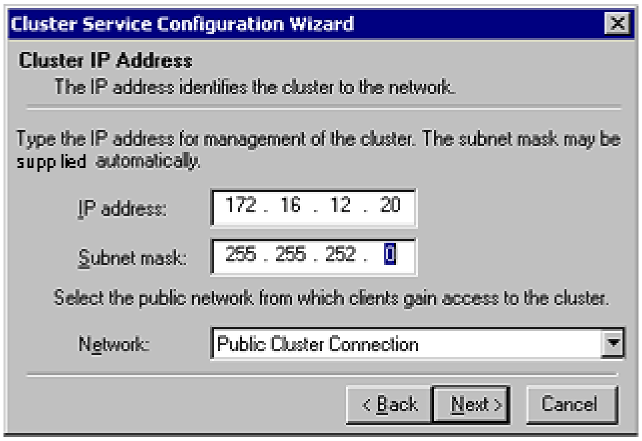

Cluster Service Configuration Wizard

Cluster IP Address

The IP address identifies the cluster to the network.

Type the IP address for management of the cluster. The subnet mask may be supp lied automatically.

IP address: 172 . 16 . 12 . 20

Subnet mask: 255 . 255 . 252 . 0

Select the public network from which clients gain access to the cluster.

Network: Public Cluster Connection

14. Enter the Cluster **IP address** and **subnet**. This must be a valid Static IP address and the proper Public Network subnet. Click **Next**.
15. Click **Finish**. At this point you should receive a message that the cluster service has been installed successfully.
16. For 2003 Server first make sure that you are part of a domain, then follow these steps:

Installing the cluster service software on Windows 2003

Note: At this time you should power on the Main node in the cluster, but leave all other nodes off.

When the Main node is powered up and you are logged into Windows with Local Administrator rights, follow these steps to install the cluster software:

1. Select **Start > Programs > Administrative Tools > Cluster Administrator**. The Cluster Administrator opens.

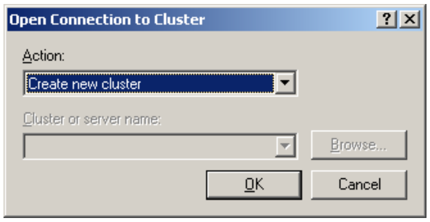

Open Connection to Cluster

Action: Create new cluster

Cluster or server name:

2. Select **Create new cluster** from the **Action** drop-down menu and click **OK**. The *New Cluster Wizard* launches.
3. Click **Next** on the Welcome screen. The *Cluster Name and Domain* screen appears.

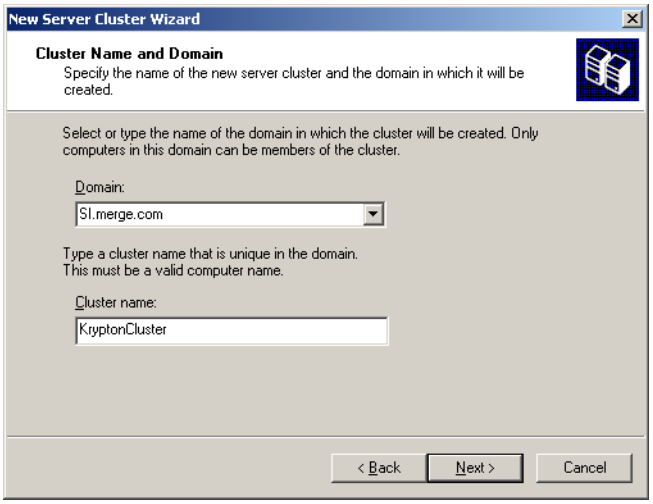

New Server Cluster Wizard

Cluster Name and Domain

Specify the name of the new server cluster and the domain in which it will be created.

Select or type the name of the domain in which the cluster will be created. Only computers in this domain can be members of the cluster.

Domain: SI.merge.com

Type a cluster name that is unique in the domain. This must be a valid computer name.

Cluster name: KryptonCluster

4. Enter your domain (if it isn't already filled in) and a cluster name and click **Next**. The *Select Computer* screen appears.

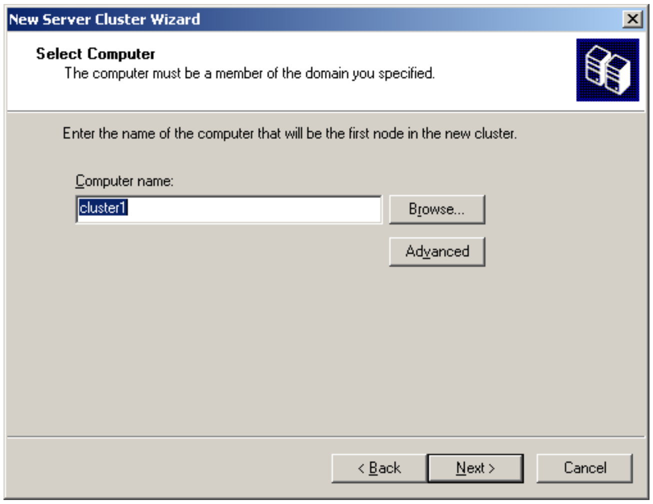

New Server Cluster Wizard

Select Computer

The computer must be a member of the domain you specified.

Enter the name of the computer that will be the first node in the new cluster.

Computer name: cluster1

5. Enter the computer name and click **Next**. The *Analyzing Configuration* screen appears.

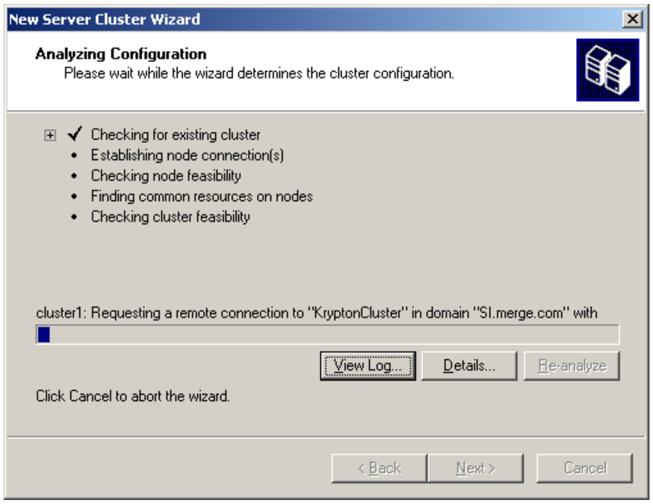

New Server Cluster Wizard

Analyzing Configuration

Please wait while the wizard determines the cluster configuration.

- ✓ Checking for existing cluster
- Establishing node connection(s)
- Checking node feasibility
- Finding common resources on nodes
- Checking cluster feasibility

cluster1: Requesting a remote connection to "KryptonCluster" in domain "SI.merge.com" with

Click Cancel to abort the wizard.

6. Wait for the wizard to finish analyzing your system configuration, then click **Next**. The *IP Address* screen appears.

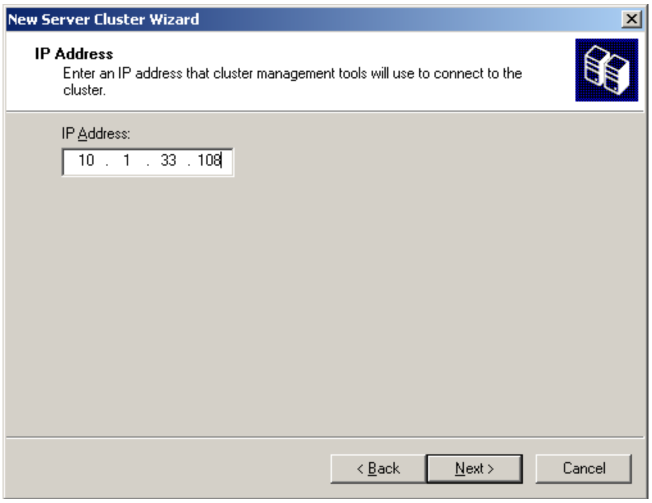

New Server Cluster Wizard

IP Address

Enter an IP address that cluster management tools will use to connect to the cluster.

IP Address: 10 . 1 . 33 . 108

7. Enter the IP address for the cluster and click **Next**. The *Cluster Service Account* screen appears.

Note: You may see an “error finding common resources” message. This occurs because the second server is turned off, and should not affect the installation.

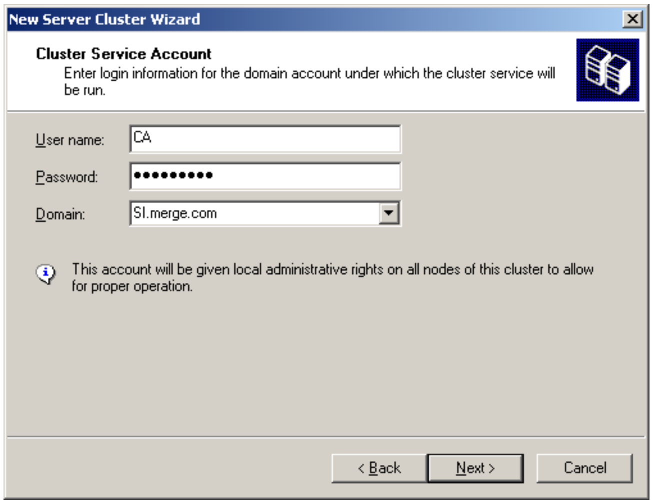

New Server Cluster Wizard

Cluster Service Account

Enter login information for the domain account under which the cluster service will be run.

User name: CA

Password: •••••••••

Domain: SI.merge.com

This account will be given local administrative rights on all nodes of this cluster to allow for proper operation.

8. Enter the domain user name and password, and select a domain. Click **Next**. The *Proposed Cluster Configuration* screen appears.

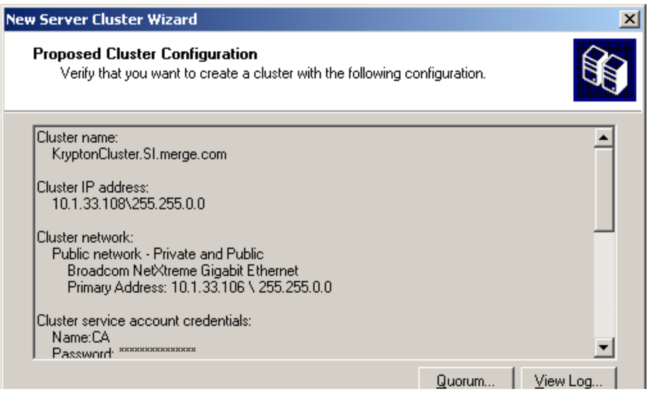

New Server Cluster Wizard

Proposed Cluster Configuration

Verify that you want to create a cluster with the following configuration.

```
Cluster name:
   KryptonCluster.SI.merge.com

Cluster IP address:
   10.1.33.108\255.255.0.0

Cluster network:
   Public network - Private and Public
      Broadcom NetXtreme Gigabit Ethernet
      Primary Address: 10.1.33.106 \ 255.255.0.0

Cluster service account credentials:
   Name:CA
   Password: ***************
```

9. Verify the information and click **Quorum**. The *Cluster Configuration Quorum* dialog box appears.

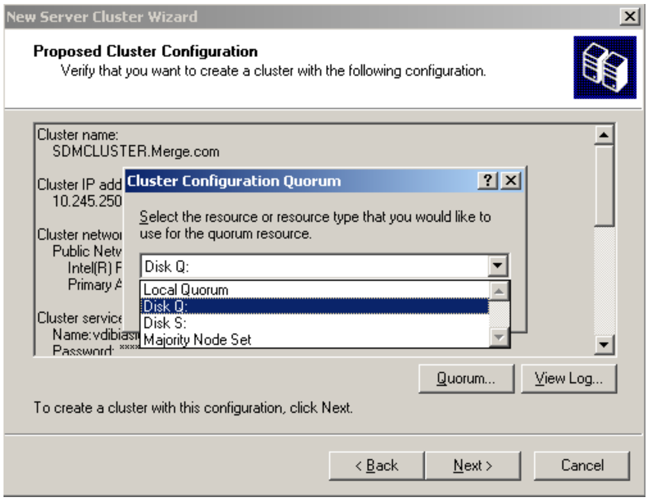

New Server Cluster Wizard

Proposed Cluster Configuration

Verify that you want to create a cluster with the following configuration.

Cluster name:
SDMCLUSTER.Merge.com

Cluster IP add
10.245.250

Cluster netwo
Public Netw
Intel(R) F
Primary A

Cluster service
Name:vdibiasi
Password: ***

Cluster Configuration Quorum

Select the resource or resource type that you would like to use for the quorum resource.

Disk Q:

To create a cluster with this configuration, click Next.

10. Verify that the correct drive has been selected for the Quorum drive and click **OK**.
11. Click **Next**. The installer will then create the cluster.

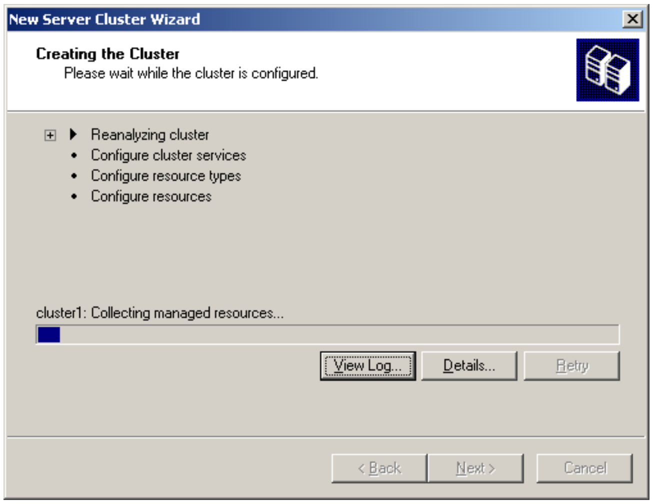

New Server Cluster Wizard

Creating the Cluster

Please wait while the cluster is configured.

- Reanalyzing cluster
- Configure cluster services
- Configure resource types
- Configure resources

cluster1: Collecting managed resources...

12. Wait for the wizard to finish creating the cluster and click **Next**.

Note: The heartbeat network info does not show up here, but it will appear after you configure the second server.

13. Click **Finish** to close the wizard.

Validating the cluster installation

To validate the cluster installation:

- Select **Start > Programs > Administrative Tools > Cluster Administrator**. If the cluster installation completed properly, you will see a cluster administrator screen similar to the picture below:

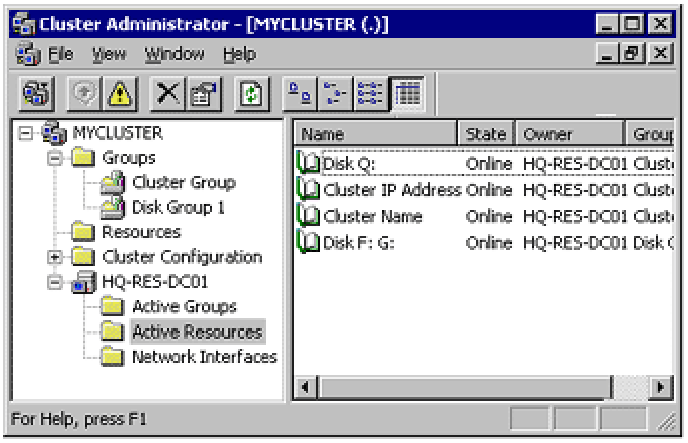

Name	State	Owner	Grou
Disk Q:	Online	HQ-RES-DC01	Clust
Cluster IP Address	Online	HQ-RES-DC01	Clust
Cluster Name	Online	HQ-RES-DC01	Clust
Disk F: G:	Online	HQ-RES-DC01	Disk (

Configuring the secondary nodes

This section describes how to add nodes to the cluster:

- if you are using one of the Windows 2000 operating systems, see "Adding a secondary node on Windows 2000" on page 88.
- if you are using one of the Windows 2003 operating systems, see "Adding a secondary node on Windows 2003" on page 89.

Adding a secondary node on Windows 2000

Note: In this section, you must leave the Main node powered on.

Power up the Secondary node. When the Secondary node is powered up and you are logged into Windows with Local Administrator rights follow these steps to attach the secondary nodes to the cluster:

1. Select **Start > Settings > Control Panel > Add/Remove Programs > Add/Remove Windows Components**.
2. Select **Cluster Service** and click **Next**. You will be prompted for the Windows Operating System installation disk.
3. Enter the appropriate path pointing to the version of Windows you are using and the **i386** directory. Click **Next**.
4. The next screen is a Microsoft disclaimer that the equipment you are using in the cluster is on the Hardware Compatibility List. Click **I Understand**, then click **Next**.
5. When you have the dialog box to create or join a cluster, select the **The second or next node in the cluster** radio button.
6. Enter the name that you chose for the cluster.

7. Leave the **Connect to cluster** check box clear. The installation will use the previous user you entered to run the Cluster Service.
8. Click **Next**.
9. Click **Finish**.

If you are installing additional nodes in the cluster, repeat the steps above for each node.

Adding a secondary node on Windows 2003

Note: In this section, you must leave the Main node powered on.

Power up the Secondary node. When the Secondary node is powered up and you are logged into Windows with Local Administrator rights follow these steps to attach the secondary nodes to the cluster:

1. Select **Start > Programs > Administrative Tools > Cluster Administrator**. The Cluster administrator launches.

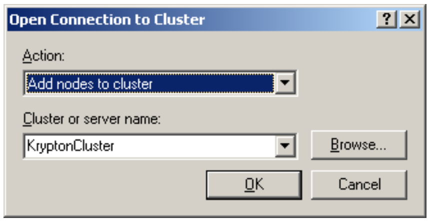

Action: Add nodes to cluster

Cluster or server name: KryptonCluster

2. Select **Add nodes to cluster** from the **Action** drop-down list.
3. Enter or select the cluster name and click **OK**. The *Add Nodes Wizard* launches.
4. Click **Next** on the Welcome screen. The *Select Computers* screen appears.

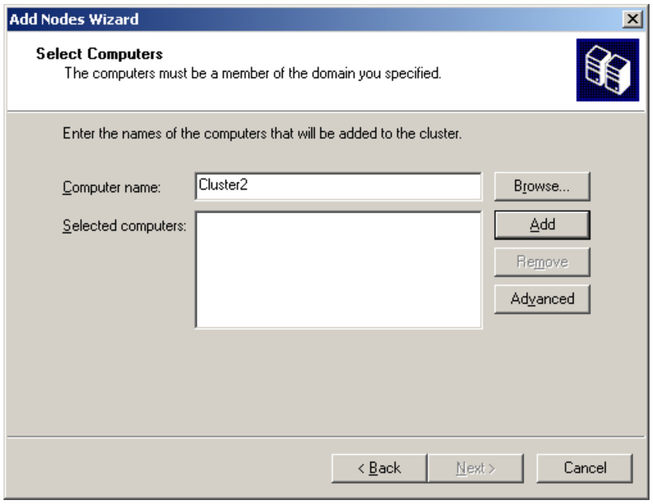

Select Computers

The computers must be a member of the domain you specified.

Enter the names of the computers that will be added to the cluster.

Computer name: Cluster2

5. Enter the machine name for the node you want to add and click **Add**. Add any other nodes that you have configured. Click **Next**. The *Analyzing Configuration* screen appears.

6. Wait for the analysis to complete and click **Next**. The *Cluster Service Account* screen appears.

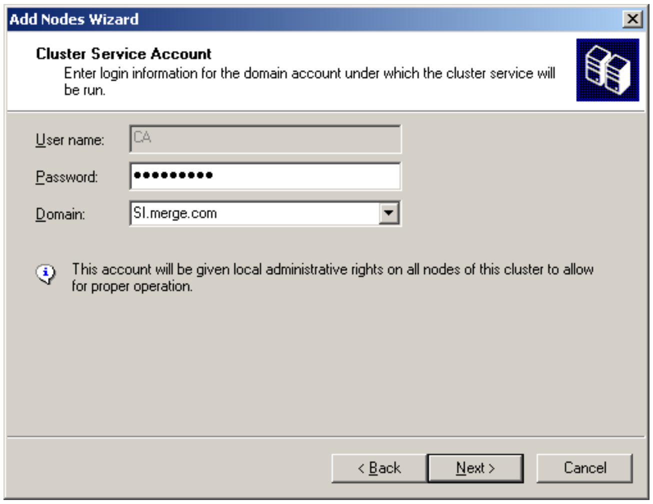

Cluster Service Account

Enter login information for the domain account under which the cluster service will be run.

User name: CA

Password: •••••••••

Domain: SI.merge.com

This account will be given local administrative rights on all nodes of this cluster to allow for proper operation.

7. Enter the user name and password for the service account, and select the domain. Click **Next**. The *Proposed Cluster Configuration* screen appears.
8. Verify all of the information and click **Next**. The wizard will start adding nodes to the cluster.
9. When the wizard has finished adding nodes, click **Next**.
10. Click **Finish** to close the wizard.

Verifying installation

1. Open the Cluster Administrator program.
2. You should now see all nodes in the cluster.
3. Right-click **Disk Group 1** and select **Move**. This option will move who is controlling these resources to the other node. As long as all devices are brought Online, then your Cluster is working.

Installing SQL Server Clustering

This section describes how to install SQL Server clustering.

Preparing for SQL Server Clustering

Before you install SQL Server make sure of the following:

- Your cluster is working properly.
- You have a Virtual SQL Server name and a static IP Address to assign to the Virtual SQL Server.
- You have configured Microsoft Distributed Transaction Coordinator (MSDTC). This can be done by running `comclust.exe` from a command prompt on each node. You can verify that MSDTC has been set up properly by opening Cluster Administrator and making sure that MSDTC is listed in the Resources folder.

Installing SQL Server 2000 Clustering

This section covers installing SQL Server 2000 Cluster.

To install SQL Server 2000 Cluster:

1. Log in to the **Main node** with a User Account that has been granted Local Administrator rights and start the SQL Sever 2000 Enterprise Edition installation.
2. Click **Next** at the *Welcome* screen. The installer will detect that this instance of SQL Server is being installed on a node that belongs to a cluster.

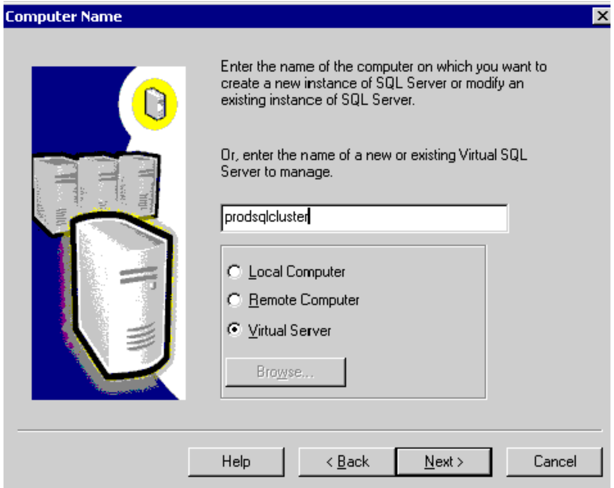

Computer Name

Enter the name of the computer on which you want to create a new instance of SQL Server or modify an existing instance of SQL Server.

Or, enter the name of a new or existing Virtual SQL Server to manage.

prodsqlcluster

○ Local Computer
○ Remote Computer
● Virtual Server

3. Enter the Virtual SQL Server name and click **Next**.
4. Enter the name and company and click **Next**.
5. The next screen is the license agreement. Read this and click **Yes**.

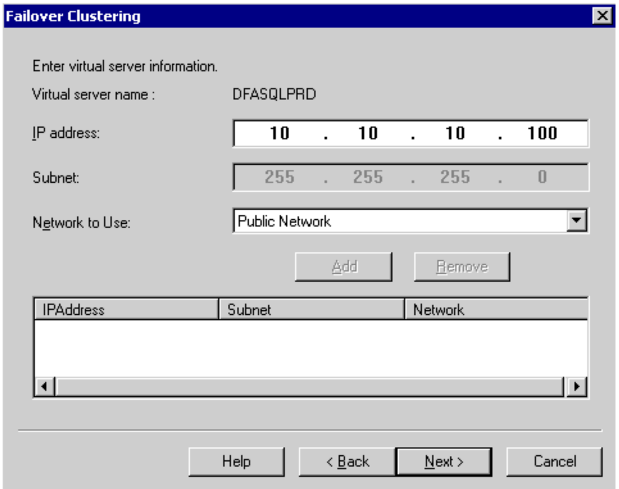

Failover Clustering

Enter virtual server information.

Virtual server name : DFASQLPRD

IP address: 10 . 10 . 10 . 100

Subnet: 255 . 255 . 255 . 0

Network to Use: Public Network

IPAddress	Subnet	Network

6. The next screen prompts you to enter the Virtual SQL Server **IP Address** and **Subnet**. Enter a valid Static IP Address and the Subnet for the Public Network. Click **Next**.

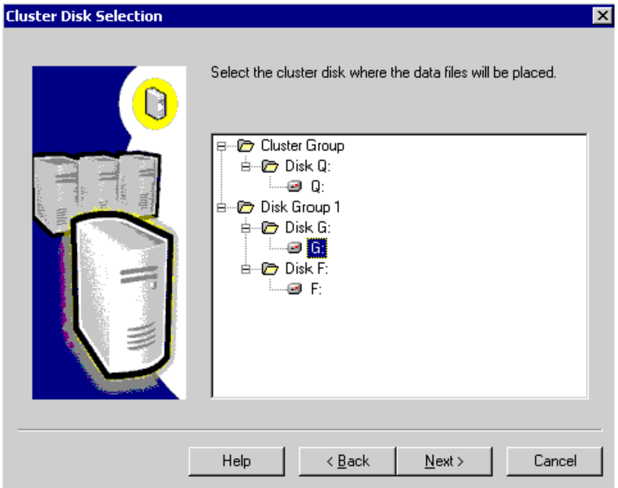

7. Indicate which shared disk to use for the SQL Data files. Click **Next**.

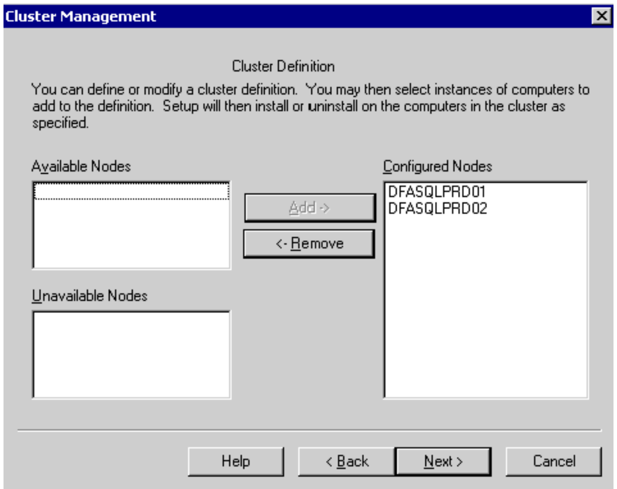
Cluster Management

Cluster Definition

You can define or modify a cluster definition. You may then select instances of computers to add to the definition. Setup will then install or uninstall on the computers in the cluster as specified.

Configured Nodes: DFASQLPRD01, DFASQLPRD02

8. Select which nodes belong to the cluster and where SQL Server should be installed. Make sure all nodes in the cluster are in the **Configured Nodes** list. Click **Next**.
9. Enter the **User name** and **password** for the Cluster Administrator. Enter the Domain Administrator Account you created in "Creating a Cluster user account" on page 78.
10. This screen allows you to set up a Default SQL Server or a named instance. We recommend that you only have a Default installation in your SQL Server Cluster. Click **Next**.
11. The next screen allows you to choose a typical installation or a custom installation. Choose the option that fits your installation and press **Next**. Depending on the type of installation you selected, you may have additional screens to complete.
12. On the *Service Accounts* screen, enter the Domain Administrator you created in "Creating a Cluster user account" on page 78. Click **Next**.
13. Enter the type of SQL authentication mode you desire and click **Next**.
14. Choose the appropriate SQL Collation and click **Next**.
15. Select the appropriate network libraries for your installation and click **Next**.

16. Enter the appropriate licensing model for your installation and click **Continue**. The installation will begin.

Important: Once the installation has completed, we recommend that you reboot all nodes in the cluster.

Installing SQL Server 2000 Service Pack

On the Main node of the cluster, follow these steps to install the SQL Server 2000 Service Pack:

1. Run **setup.bat** from the SQL Server 2000 Service Pack folder.
2. The *Welcome* screen appears for the Service Pack. Click **Next**.

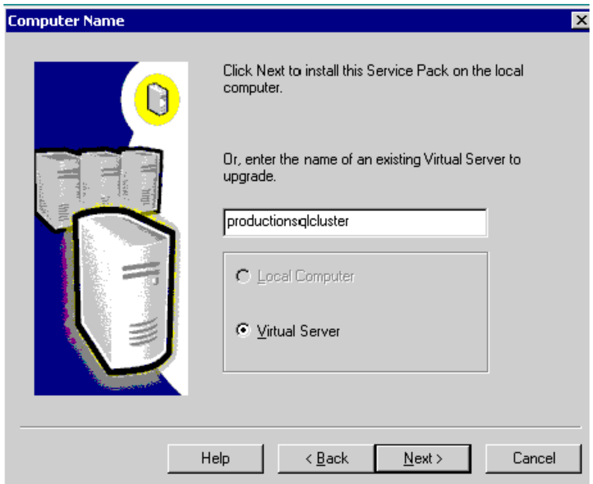

Computer Name

Click Next to install this Service Pack on the local computer.

Or, enter the name of an existing Virtual Server to upgrade.

productionsqlcluster

○ Local Computer

● Virtual Server

3. Enter the name of the virtual SQL Server and select the **Virtual Server** radio button. Click **Next**.
4. Select the Authentication mode to install the service pack with.
5. The installation will verify the user name and password.
6. The next screen will prompt with a few options. Make sure the required options are checked and click **Next**.
7. You might receive a message to send fatal errors to Microsoft. Select an option and click **Next**.
8. Enter the administrator account for the Cluster. Enter the user name and password you created in "Creating a Cluster user account" on page 78 and click **Next**.
9. The installer will copy files. When the installation is complete, click **Finish**.

Verifying the SQL Server Cluster installation

You can verify that the SQL Server Cluster installation was successful by opening the Cluster Administrator program. Highlight the node Disk Group 1. On the right pane you should see a series of SQL Server resources and they should all be online.

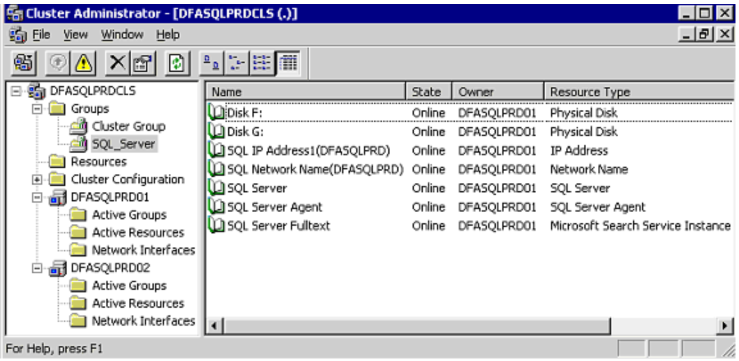

Name	State	Owner	Resource Type
Disk F:	Online	DFASQLPRD01	Physical Disk
Disk G:	Online	DFASQLPRD01	Physical Disk
SQL IP Address1(DFASQLPRD)	Online	DFASQLPRD01	IP Address
SQL Network Name(DFASQLPRD)	Online	DFASQLPRD01	Network Name
SQL Server	Online	DFASQLPRD01	SQL Server
SQL Server Agent	Online	DFASQLPRD01	SQL Server Agent
SQL Server Fulltext	Online	DFASQLPRD01	Microsoft Search Service Instance

You can also rename the Disk Group 1 to something a little more meaningful. In the example above the Disk Group 1 folder has been renamed to SQL_Server.

Troubleshooting

If, for any reason, your SQL Server Cluster has not installed properly you can examine the cluster log file to see if there is any detailed information on why the cluster service and SQL Server cluster is failing. The Cluster log file can be located in **%root%\WINNT\Cluster\cluster.log**.

You can also examine the Event Viewer to see if there is extended error information for the Cluster service.

Setting up Fusion PACS for use with the cluster

This section describes how to set up Fusion PACS SDM for use with a SQL Server Cluster.

To set up Fusion PACS SDM with SQL Server Clustering:

1. Perform a full installation of the PACS (web, services and database) on the server where you are placing your services and web.
2. Edit the following parameters in the **web.config** file to point at your virtual SQL server's IP address (edit the SA and password values as needed):
   - FusionServerConnectionString
   - FusionServerUserManagerConnectionSting
3. In Windows, run **regedit**.
4. Navigate to **HKEY_LOCAL_MACHINE/Software/Merge eFilm/Fusion Server/ Connection** and modify the IP, SA and password values as needed to point them at your virtual SQL server's IP.
5. Close regedit.
6. Copy the following scripts to the main node of the cluster:
   - CreateFusionDB.sql
   - CreateFusionTables.sql
   - CreateFusionSP.sql
   - InsertFusionData.sql
   - CreateFusionUserManDB.sql

- CreateFusionUserManTables.sql
- CreateFusionUserManSP.sql
- FusionUserManInsertData.sql

7. On the main node of the database cluster, edit the **CreateFusionDB.sql** and **CreateFusionUserMan.sql** scripts. Modify the path to install to the root of your database drive on the cluster.
8. Using SQL Query Analyzer, run the scripts in the following order:

   a) CreateFusionDB.sql

   b) CreateFusionTables.sql

   c) CreateFusionSP.sql

   d) InsertFusionData.sql

   e) CreateFusionUserManDB.sql

   f) CreateFusionUserManTables.sql

   g) CreateFusionUserManSP.sql

   h) FusionUserManInsertData.sql
9. On the Services machine, it may be necessary to register the ASP.Net component using the following command:

   ```
   ASPNET_REGIIS.exe -I
   ```

   Note: This program can be found in the Windows\Microsoft.net\Framework\v1.1.4.2322 directory.

Service Notes B

This section collects various service bulletins issued since the release of 1.27 that are still relevant to the current release. The following topics are discussed:

- Verifying database integrity

Verifying database integrity

Some customers have expressed concern that data may be lost at their sites; other sites have lost databases completely. This document describes how to test the database at a client site to ensure data integrity.

Engineering has created a tool called CheckFileSystem that allows us to test a variety of critical conditions with the database and image data. This tool includes a switch to rebuild a lost database without re-running the data through the dicom/storage service. The utility steps through a number of checks to make sure the system is handling image data properly:

- Reads in all DCM files to make sure they are readable
- Checks for proper existence of wavelet files (single & multiframe)
- Checks that wavelet files are not zero length
- Checks compressed files are missing their pixel data
- Verifies that the Study & Series Instance UIDs in file match directory names
- Verifies that the Series Instance UID from directory name exists in the database
- Verifies that the Study Instance UID in the database for Series matches the directory name
- Verifies that the Number of Series Related Instances in DB matches image count.
- Identifies series & study level directories with no files
- Looks for stray .tmp and .lock files in directories
- Checks whether archive.tmp folders found have entries in Library_ArchivalQueue or Location
- Counts the number of header files

This utility is included in the installation and resides in the Fusion PACS installation directory (by default, C:\Program Files\Merge eFilm\Fusion Server). We highly recommend running this tool at every client site.

To verify database integrity:

1. Download CheckFileSystem.exe from:

   `ftp://ftp3.merge.com/Outgoing/Fusion/1.5/1.5.0/tools/`

Note: This version of the utility can only be used with Fusion PACS 1.5; the CFS utility from release 1.4.2 cannot be used to check a release 1.5 file system.

2. Copy the file into the directory where the SDM is installed.
3. Open a command window and navigate to the SDM installation directory.
4. Type the following command and press **Enter** to list the options available for the tool:

   ```
   CheckFileSystem.exe -h
   ```
5. Run the tool as appropriate. The utility creates an output file: **<FusionInstallDirectory>/log/<CurrentDate>/XXX_checkfilesystem.log** (the XXX is the PUID of the running CFS process).

Note: This file is created every time you run CheckFileSystem, regardless of the options chosen. If you use the -o option to create a .cvs output file, it is important to know that the .csv file will only get created the utility detects an error; otherwise only the .log file will be created.

Upgrading From 1.4.2

C

This section describes the procedures for upgrading to Fusion PACS 1.5 from an existing 1.4.2 installation.

This section also describes:

- How to upgrade the SDM (see "Upgrading the SDM" on page 99)
- How to uninstall the PACS following a failed installation (see "Uninstalling the PACS" on page 102)
- Common problems you may encounter during an upgrade (see "Common problems" on page 103)
- The database options that are available during an upgrade or installation and how they work, including the scripts that are run for each option (see "Database options during installations" on page 104)

Upgrading the SDM

This section describes the steps required to upgrade from Fusion PACS Storage and Distribution Manager 1.4.2 to Fusion PACS Storage and Distribution Manager 1.5.

To upgrade the SDM:

1. Make sure your system meets the minimum requirements for an upgrade (see "Minimum version required to upgrade" on page 99)
2. Back up the databases (see "Backing up the SDM databases" on page 99).
3. Uninstall Fusion PACS 1.4.2 (see "Uninstalling the existing version" on page 100).
4. Run the SDM installation (see "Running the upgrade" on page 101).
5. If the site is using database replication, follow the procedure described in "Upgrading the Publisher and Subscriber databases" on page 101 to upgrade the databases to use SQL log shipping.

Minimum version required to upgrade

The minimum required FUSION Server product needed in order to perform an upgrade is 1.4.2. If the customer is using an earlier version of FUSION Server, you will have to upgrade that customer to version 1.4.2 before upgrading to version 1.5.

Backing up the SDM databases

Before doing an upgrade from FUSION Server 1.4.2 to Fusion PACS 1.5, always perform a complete backup of the **FusionServer** and **FusionServerUserManager** databases. Verify that the backups were properly written to the hard drive up by physically viewing them. **Do not proceed** with an installation unless this step is completed successfully.

To back up the databases:

1. Open Enterprise Manager.
2. Locate the **FusionServer** database in the SQL tree view.
3. Right-click the **FusionServer** database icon and select **All Tasks > Backup Database...** The *SQL Server Backup* dialog box appears.
4. On the *General* tab, verify that the Database is **FusionServer**.
5. In the *Backup* section, verify that the radio button **Database - complete** is selected.
   a) If no backup location is defined for the database, click **Add**. The *Select Backup Database* dialog box appears.
   b) Enter a file name and location for the database and click **OK**. The dialog box closes.
6. Click **OK** to start the backup.

   Note: If the backup fails, it's because the FUSION OR is locking the tables. In that case, stop the FUSION Services to complete the backup.

7. Repeat the above steps for the **FusionServerUserManager** database.

Uninstalling the existing version

In previous versions of the FUSION SDM Server software, InstallShield would automatically uninstall an old version of FUSION SDM when detected. This automatic uninstall of a previous version, while still attempted and allowed when upgrading, is a process that we prefer you perform manually.

We recommend the following steps:

1. Perform a proper manual shutdown (with confirmation) of all FUSION services on all FUSION SDM servers in the cluster prior to the upgrade.

   Note: Please consider that other FUSION SDM Servers in a cluster may be running. Stop ALL FUSION services (such as DISK Service, Storage Service, etc.) on all FUSION SDM Servers before updating the main FUSION Server database. You may want to halt your web server if actively being used by remote web clients.

2. The database server (if in a cluster) should be upgraded before any other servers in the cluster.
3. After the main database (and perhaps the subscriber database, if it exists) has been upgraded, the remaining servers in the cluster may be upgraded.

A given FUSION SDM server may be running the database, the runtime services, the web server or any combination thereof. Regardless of the combination, we recommend a manual uninstall of the previous version (via the Control Panel), before installing the new version.

Running the upgrade

An upgrade follows the same procedure as a normal Fusion PACS installation (as described in Chapter 1, "Installing Fusion PACS"), with one important difference: when setting up the databases, you must select **Update the Fusion Databases** as shown below.

WARNING! Failure to set this option properly can erase the FUSION databases.

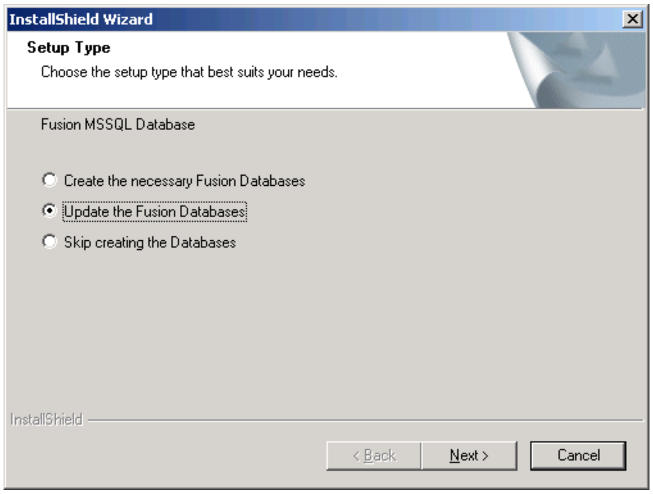

InstallShield Wizard
Setup Type
Choose the setup type that best suits your needs.
Fusion MSSQL Database
○ Create the necessary Fusion Databases
● Update the Fusion Databases
○ Skip creating the Databases
InstallShield
< Back
Next >
Cancel

Also note that when upgrading a previous installation of the FUSION Server to FUSION 1.4.2, the installation scripts called by InstallShield will add SQL Server indices to a number of existing tables in both the FusionServer and FusionServerUserManager databases. These indices significantly improve the response time of database queries and data inserts in a number of key areas.

The generation of this table indexing data occurs as the upgrade scripts run and can take a significant amount of time. Sites with many hundreds of thousands of entries in the SeriesLocation table, the archival queue table, and/or with many hundreds of users may take as long as an hour for this install script to complete.

Important: Please allow for up to an hour when scheduling a database upgrade and be **extremely patient** while the InstallShield code performs the database upgrade.

These long delays do not occur when installing a new FUSION Server, as there is no existing database to upgrade.

Upgrading the Publisher and Subscriber databases

If the site is set up to use database replication, you cannot simply upgrade both the Publisher and Subscriber database servers to 1.5, since database replication has been replaced with SQL log shipping. Instead, you must turn off replication and set up SQL log shipping using the Publisher as the Primary and the Subscriber as the backup server.

To perform an upgrade in a replication scenario:

1. **Drop both subscriptions from the Replication Server.** Using Microsoft SQL Server Enterprise Manager, perform the following on the failover server:

a) In the tree pane on the left, expand the **Replication** heading.

b) Select the **Subscriptions** element beneath Replication

c) In the view pane select a subscription and right-click.

d) Select **Stop Synchronizing**.

e) Right click the subscription again and select **Delete**, then click **Yes**.

Repeat the above steps for the other subscription.

2. **Disable distribution on the Replication (failover) Server.** Using Microsoft SQL Server Enterprise Manager, perform the following on the failover server:

   a) Right-click the server registration for the failover server.

   b) Select **Properties**.

   c) Select the **Replication** tab of the dialog box.

   d) In the *Publishing and distribution group* box, click **Disable**; this launches a wizard to guide you through the process of disabling distribution.

   e) When the wizard appears, click **Next**.

   f) Under *Disable Distribution*, select **Yes** (disable publishing), then click **Next**.

   g) Click **Next** to confirm, then click **Finish**.

3. **Delete Replication jobs/agents.** Using Microsoft SQL Server Enterprise Manager, perform the following on the failover server:

   a) In the tree pane on the left expand the **Management** heading.

   b) Expand the **SQL Server Agent** element beneath Management.

   c) Select the **Jobs** element beneath SQL Server Agent.

   d) Right-click **Jobs** and select **Refresh**.

   e) In the view pane select all jobs whose category begins with **REPL**.

   f) Delete the selected jobs.

4. Uninstall **Fusion Server** from the Subscriber machine.

5. Manually delete the **FusionServer** and **FusionServerUserManager** databases from the Subscriber machine.

6. On the **Publisher** machine, install FUSION Server 1.5 normally.

7. Follow the instructions in the section "Using SQL log shipping" on page 24 to set up SQL log shipping.

Uninstalling the PACS

You may need to uninstall the upgrade and start again with a fresh install of 1.4.2.

To perform a complete uninstallation of 1.5:

1. Use **Add/Remove Programs** to uninstall Fusion PACS 1.5.
2. Use *Enterprise Manager* to remove the **FusionServer** and **FusionServerUserManager** databases.
3. Reinstall **Fusion Server 1.4.2**.
4. Restore the database (you DID back up the database, right?).
5. Retry the Fusion PACS 1.5 installation.

Common problems

This section describes common problems that you may encounter when upgrading the system.

Note: This section only includes problems that were known to PI at the time of release. For the latest news, check DocuShare for any post-release service bulletins.

This section describes the following situation:

- Services missing after an upgrade (see "Services missing after an upgrade" on page 103)

Services missing after an upgrade

On a rare instance one or two of the NT services (DICOM, Disk, Storage, etc.) will not be removed during the uninstall of FUSION Server 1.4.2. This is caused when an NT service is locked and the installer is trying to remove it. If this happens, the symptoms experienced will be missing NT services following an upgrade installation of FUSION Server 1.5. The installer will also report that it was unable to stop or remove the service in question while doing the FUSION Server 1.4.2 uninstall. This does not mean that you must do a new install. This can be remedied by manually installing the services as described below.

To manually install the NT services:

1. Once the FUSION Server 1.5 install is complete, open the NT Services application and verify that the following services are present:
   - Dicom
   - Disk
   - IC
   - Media
   - Move
   - Replication
   - Storage Sync Services
2. Open a command window.
3. For each missing service, run the appropriate manual installation command. This example shows the commands for the Dicom, Disk, Storage, and Move services:

   ```
   "c:\Program Files\Merge eFilm\Fusion Server\dicomservice.exe" -Service
   "c:\Program Files\Merge eFilm\Fusion Server\diskservice.exe" -Service
   "c:\Program Files\Merge eFilm\Fusion Server\storageservice.exe"
   -Service
   "c:\Program Files\Merge eFilm\Fusion Server\moveservice.exe" -Service
   ```
4. Before running the commands, navigate to the directory where FUSION Server is installed—that way, you can omit the directory path when typing each command.
5. Once the command is run, locate the newly-created service in the NT services and double click it. The *Properties* dialog box for that service appears.

   On the **General** tab:

   a) Add the word **Fusion** in front of the Display Name Value.

b) Select the appropriate **Startup Type**.

On the **Log On** tab:

- If this account is set up to run under a user context ,then select the This account radio button, select an account, and complete the Password and Confirm password fields. Mimic the other Services that were installed correctly.

6. Click **OK**.

Database options during installations

The following is an explanation of the various options presented as answers to the database question asked during an installation of FUSION Server 1.5. The correct response to the database depends on the scenario of the install. If the wrong option is selected, the installation will need to be redone. If the database is not backed up (but you did back up the database, didn't you?), the data could be gone forever.

There are three options, explained below:

- Create New
- Update Database
- Skip Database

Create New

This is for installing brand new installations of FUSION Server 1.5. If this option is selected, it creates new FusionServer and FusionServerUserManager databases in MSSQL. If a previous version of FUSION Server was present, this will destroy the old databases and replace them with two new empty databases .

Here's what the SQL scripts do when you select "Create New" (the SQL status can be found in <FUSION install directory>/Script/Logs):

1. Create the FusionServer database.
2. Create the FusionServerUserManager database.
3. Create the FusionServer tables.
4. Create the FusionServerUserManager tables.
5. Populate the FusionServer tables.
6. Populate the FusionServerUserManager tables.
7. Create the FusionServer Stored Procedures.
8. Create the FusionServerUserManager Stored Procedures.

Update Database

This is for installing from one released version of FUSION to another released version of FUSION. In the case of Fusion PACS 1.5, Service should select Update Database when going from Fusion PACS 14 to Fusion PACS 1.5. Update should only be run once when going from release to release. Do not try to re-run an install and use the update option a second time.

If you run Update to get a database to the 1.5 version of FUSION Server and then uninstall FUSION Server 1.5, the database is still present in MSSQL and remains at version 1.5. If you use the update option again, the data will be wrecked and the database may need to be reinstalled by using the Create New database option described above and then restoring from a backup.

Here's what the SQL scripts do when you select "Update Database" (the SQL status can be found in <FUSION install directory>/Script/Logs):

1. Update Script for fixing the schema and populating additional data to FusionServer and FusionServerUserManager.
2. Create the FusionServer Stored Procedures.
3. Create the FusionServerUserManager Stored Procedures.

Skip Database

This is a seldom-used option that can be used to keep the installer from doing any database manipulation. This could be used if an installation updated the databases to the Fusion PACS 1.5 version, and then the FUSION software was uninstalled via *Add/Remove Programs*. The installer could reinstall the 1.5 PACS and select Skip Database. This would keep the database from re-running any update scripts that would change schema formats. However, the Stored Procedures will be updated to the FUSION 1.5 release to match the binaries installed.

Note: Do **not** use this option if you are upgrading FUSION Server 1.4.2 databases to the latest release. This could cause problems when stored procedures are called during the normal operation of FUSION Server, because the version of the database and binaries/ stored procedures would be mismatched.

Here's what the SQL scripts do when you select "Skip Database" (the SQL status can be found in <FUSION install directory>/Script/Logs):

1. Create the FusionServer Stored Procedures.
2. Create the FusionServerUserManager Stored Procedures.

Contacting Merge Healthcare D

If the procedures in this manual do not help you solve the problem, or the symptoms you are experiencing do not appear in this manual, contact Merge Healthcare for assistance.

USA

6737 W. Washington Street, Suite 2250
Milwaukee, WI 53214-5650, USA.
Tel: 414-977-4000
Toll Free: 1-877-741-5369
FAX: 414-977-4200
E-mail: support@merge.com

Canada

6509 Airport Road
Mississauga, Ontario CANADA
L4V 1S7
Tel: 1-416-672-9425
Toll Free: 1-877-741-5369
FAX: 1-416-672-2307
E-mail: support@merge.com

Europe

Spegelt 34
5674 CD Nuenen
The Netherlands
Tel: (31) (40) 2990773
Fax: (31) (40) 2906615
E-mail: service_europe@merg.com

World Wide Web

www.merge.com

Before you call

Before calling Merge Healthcare for assistance, please prepare the following information:

- Site name and location
- System Administrator’s name and contact information
- Detailed description of the problem
- Detailed description of troubleshooting attempts

Index

I