

System Administrator User Guide

- [Preface](#)
 - [Getting Help](#)
 - [Conventions](#)
- [Introduction](#)
 - [About This Guide](#)
 - [DFdiscover System Overview](#)
 - [Software Version](#)
 - [Client/Server Terminology](#)
 - [Daemons and Other Background Processes](#)
 - [Master Daemon](#)
 - [Study Server](#)
 - [EDC server](#)
 - [Inbound Server](#)
 - [Outbound Server](#)
 - [DFdiscover System Limits](#)
- [Security](#)
 - [Secure Communication](#)
 - [User Authentication](#)
 - [Two-Factor Authentication](#)
 - [Microsoft Entra ID Single Sign-On \(Formerly Azure AD\)](#)
 - [Restrict Login to Microsoft Entra ID SSO Only](#)
 - [Azure SMS Messaging](#)
 - [Server Filesystem Permissions](#)
 - [PDF Encryption Password](#)
 - [DFdiscover System-Level Permissions](#)
 - [DFdiscover Study-Level Permissions](#)
 - [Impact of Database Access Restrictions on Standard Reports](#)
 - [Impact of Database Access Restrictions on Edit Checks](#)
 - [ePRO Scheduling Permissions](#)
 - [Study Filesystem Permissions](#)
- [Getting Started with DFadmin](#)
 - [Starting DFadmin](#)
 - [DFdiscover System Administration Tasks](#)
- [License](#)
 - [Introduction](#)
 - [DFadmin's License Tab](#)
 - [Entering and Reviewing License Information](#)
 - [Password Restricted Values](#)
 - [Password Independent Values](#)
 - [Applying Changes](#)
 - [License Expiry](#)
 - [Updating License Information without DFadmin or DFserveradmin](#)
 - [Reviewing License Usage Statistics](#)
- [Master](#)
 - [Starting DFdiscover](#)
 - [Stopping DFdiscover](#)
 - [Forcing DFdiscover to shutdown](#)
 - [Configuring the DFdiscover master](#)
- [Studies](#)
 - [Study Numbers](#)
 - [Study Directories](#)
 - [The Studies Tab](#)
 - [Adding a New Study](#)
 - [Deleting a Study](#)
 - [Modifying an Existing Study](#)
 - [Watermarks](#)
 - [Study History](#)
- [Roles](#)
 - [The Roles Tab](#)
 - [Adding a New Role](#)
 - [Copying and Pasting a Role](#)
 - [Deleting a Role](#)
 - [Tools & Reports](#)
 - [DFexplore Views](#)
 - [DFexplore Print/Save](#)
 - [DFexplore Miscellaneous](#)
 - [API Clients](#)
 - [Server](#)
 - [Setup](#)
 - [Auto Logout](#)
 - [Database Permissions](#)

- [Adding Database Permission](#)
 - [Deleting Database Permission](#)
 - [Modifying Existing Study Permissions](#)
 - [Restrictions](#)
 - [Role History](#)
 - [Assign Roles to Users](#)
 - [Import](#)
 - [Export](#)
- [Users](#)
 - [The Users Tab](#)
 - [User Profile](#)
 - [Adding a New User](#)
 - [Duplicating an Existing User](#)
 - [Deleting an Existing User](#)
 - [Filtering Users](#)
 - [Adding New Users in Bulk](#)
 - [Permissions](#)
 - [Adding Study Permissions](#)
 - [Deleting Study Permissions](#)
 - [Modifying Study Permissions](#)
 - [ePRO Admin Permissions](#)
 - [User History](#)
 - [User and Session Settings](#)
 - [Sharing Settings File\(s\)](#)
 - [Local Cache](#)
 - [Import](#)
 - [Export](#)
- [Status](#)
 - [Changing Study Status](#)
- [Traffic](#)
 - [Incoming Documents](#)
 - [Processing of an incoming document](#)
 - [Processing of an incoming email](#)
 - [Outgoing Documents](#)
 - [Inbound](#)
 - [Modifying Inbound Configuration](#)
 - [Enable/Disable Inbound](#)
 - [Enable/Disable High Definition \(HD\) Imaging](#)
 - [Outbound Service](#)
 - [Modifying Outbound Configuration](#)
 - [Start/Stop Outbound](#)
 - [Auto Reply](#)
 - [Registered Sites Database](#)
 - [Auto reply receipt for registered sites](#)
- [DFserveradmin](#)
 - [Master](#)
 - [Request an SSL certificate signing](#)
 - [Installing a Signed Certificate](#)
 - [License](#)
 - [Study Spaces](#)
 - [Study Diagnostics](#)
 - [Study Permissions](#)
- [Periodic Maintenance](#)
 - [Disk Maintenance](#)
 - [Monitoring and planning future disk needs](#)
 - [Disaster Planning and backup/recovery procedures](#)
 - [Replicating your server to another location](#)
 - [Disk backup and recovery procedures](#)
 - [Archive File Maintenance](#)
 - [Study Maintenance](#)
 - [Going live with a **DFdiscover** study](#)
 - [Monitoring study directory permissions](#)
 - [Cleaning the study *work* directory](#)
 - [Cleaning the study Query Reports directory](#)
 - [Procedures for archiving a **DFdiscover** study](#)
 - [Retrieving lost CRF images](#)
 - [DFdiscover Study Consistency](#)
 - [DF_ICrecords](#)
 - [DF_ICimages](#)
 - [DF_ICqcs](#)
 - [System Maintenance](#)
 - [Certificate Info](#)
 - [Truncating **DFdiscover** log files](#)
 - [Files to not prune](#)
 - [Truncating **HylaFAX** log files](#)
 - [Conclusion](#)
- [Troubleshooting](#)

- [System Error Log](#)
 - [Message Format](#)
- [Login Error Messages](#)
- [Locks](#)
 - [/opt/dfdiscover/work/DFmaster.rpcd.lock](#)
- [Diagnosing a study server](#)
- [Troubleshooting Techniques](#)
 - [Check for DFdiscover or system error messages](#)
 - [Permissions: does the problem occur as user datafax?](#)
 - [Permissions: are the DFdiscover application permissions correct?](#)
 - [Which step of a shell-script is failing?](#)
- [Other troubleshooting areas](#)
 - [Database server process cannot be started](#)
 - [CRF images are missing, or partially processed documents are found in an incoming daemon work directory](#)
- [Recovering from a Power Failure](#)
- [Recovering from a full Filesystem](#)
 - [Interaction with HylaFAX](#)
- [DFdiscover System Administration Tool](#)
 - [Auto Logout](#)
 - [DFdiscover](#)
- [DFdiscover System Files](#)
 - [DFdiscover system directories](#)
 - [Format used to describe files](#)
 - [DFdiscover system configuration files](#)
 - [DFdiscover system work files](#)
- [Copyrights](#)
 - [External Software Copyrights](#)
 - [DCMTK software package](#)
 - [Jansson](#)
 - [Mimencode](#)
 - [RSA Data Security, Inc., MD5 message-digest algorithm](#)
 - [mpack/munpack](#)
 - [TIFF](#)
 - [PostgreSQL](#)
 - [OpenSSL License](#)
 - [Original SLEay License](#)
 - [gawk](#)
 - [Ghostscript](#)
 - [MariaDB and FreeTDS](#)
 - [QtAV](#)
 - [FFmpeg](#)
 - [c3.js](#)
 - [d3.js](#)
 - [jwt-cpp](#)
 - [QXlsx](#)

Preface

DFdiscover Release 5.11.0

All rights reserved. No part of this publication may be re-transmitted in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of DF/Net Research, Inc. Permission is granted for internal re-distribution of this publication by the license holder and their employees for internal use only, provided that the copyright notices and this permission notice appear in all copies.

The information in this document is furnished for informational use only and is subject to change without notice. DF/Net Research, Inc. assumes no responsibility or liability for any errors or inaccuracies in this document or for any omissions from it.

All products or services mentioned in this document are covered by the trademarks, service marks, or product names as designated by the companies who market those products.

Google Play and the Google Play logo are trademarks of Google LLC. Android is a trademark of Google LLC.

App Store is a trademark of Apple Inc.

June 01 2026

Copyright © 2026 DF/Net Research, Inc.

Getting Help

For software support, Please contact the DFdiscover team:

- via email, help@dfnetresearch.com.
- Visit our website, <https://www.dfnetresearch.com>.

Conventions

A number of conventions have been used throughout this document.

Any freestanding sections of code are generally shown like this:

```
# this is example code
code = code + overhead;
```

If a line starts with # or %, this character denotes the system prompt and is not typed by the user.

Text may also have several styles:

- Emphasized words are shown as follows: ***emphasized*** words.
- Filenames appear in the text like so: dummy.c.
- Code, constants, and literals in the text appear like so: main.c.
- Variable names appear in the text like so: nBytes.
- Text on user interface labels or menus is shown as: **Printer name**, while buttons in user interfaces are shown as `Button`.
- Menus and menu items are shown as: `File` > `Exit`.

Introduction

About This Guide

This guide describes all aspects of **DFdiscover** system administration and is written for **DFdiscover** system administrators. Most tasks can be performed through the **DFdiscover** system application, **DFadmin**, but some require typing UNIX commands as either *datafax* or *root*. Some familiarity with UNIX and computer system administration is assumed.

NOTE: This guide does not describe **DFdiscover** installation or configuration procedures. Those topics are covered in [Software Installation Guide](#). In this guide, it is assumed that the installation and initial configuration of **DFdiscover** has been completed.

This chapter begins with a system level overview, meant to give an appreciation of how **DFdiscover** works. [Security](#) presents an overview of security topics related to **DFdiscover**.

[Getting Started with DFadmin](#) through [Traffic](#) describe how to use the **DFdiscover** system administration application, **DFadmin**. This application is used to configure new studies, roles and users, manage your **DFdiscover** license, start and stop **DFdiscover** daemons, and several other useful things.

Several administration functions are available when **DFdiscover** is not running. [DFserveradmin](#) describes **DFserveradmin** which provides these functions.

[Periodic Maintenance](#) describes recommended periodic maintenance procedures, which are very important to the health of your **DFdiscover** system.

[Troubleshooting](#) provides troubleshooting and system crash recovery assistance.

[DFdiscover System Files](#) describes the **DFdiscover** file system and its important directories and files.

DFdiscover System Overview

The goal in this chapter is to give you a general idea of how **DFdiscover** works at the server level, a level not typically seen by the end user. We describe the main system components and how they interact. Although we sometimes refer to file and command names, the purpose of this chapter is not to describe how to perform system administration tasks. That will come in the chapters to follow. This introduction will help you to see how the various system administration tasks fit into the bigger **DFdiscover** system picture.

Software Version

In **DFdiscover**, a version number comprises 3 components: a major version number (X), a minor version number (Y) and a patch number (Z). Every client and server application is identified with a version number of the form X.Y.Z.

A major release of the software increments the major version number (X+1) and resets the minor version and patch numbers back to 0. Major releases can include new functionality, protocols and data structures that impact compatibility with previous releases.

Minor releases can include bug fixes and functionality updates/improvements that work on top of the existing protocols and data structures. New functionality may also be introduced so long as it does not require changes to the software infrastructure on the server. Patches may be released for individual applications to fix urgent bugs. Patches do not introduce any new functionality.

Every **DFdiscover** application has a X.Y.Z version number. Generally speaking, all of the applications have the same number. When a client application connects to a server application they compare version numbers. The major version numbers must always match. If they do not, the connection immediately fails and the client reports an error. If the minor version numbers do not match, the success of the connection is controlled by the **Version Strict** settings defined for **DFdiscover** and possibly also at the study level. Those settings are described in [Version Strict \(Master\)](#) and [Version Strict \(Studies\)](#).

Client/Server Terminology

DFdiscover is a client/server application which means that the results supplied by the software are not generated by one monolithic application, but rather by multiple, co-operating processes. Some of these co-operating processes are the client applications that users interact with, while other processes are non-interactive, and computationally oriented. The former is called a client and the latter is called a server. In general, a client is any application that a user can invoke from a command-line or from a window system menu. In **DFdiscover**, typical clients include **DFexplore**, **DFweb**, **DFcollect**, **DFengage**, **DFsetup**, and **DFadmin**, as well as command-line clients such as **DFexport**.

Server applications are more difficult to describe and categorize because there are many different types of servers. In general, a server is an application that a client application can ask to perform an action on its behalf. Your UNIX workstation, as configured by the operating system, may already be a login server (managing login requests from users), a boot server (providing system files to other diskless computers so that they can boot), and a file server (providing disk file access to other computers).

In **DFdiscover**, there are study database servers (handling all requests to a study database), an EDC server (which handles secure, encrypted requests from data collection tools (**DFexplore/DFweb/DFcollect/DFengage**) and **DFadmin** clients), an incoming document server (routing all incoming documents to their correct location), and an outgoing document server (handling all requests to send out documents).

Daemons and Other Background Processes

While most users will only be aware of the individual applications they use, these tools are really just clients which request services from one or more of the main **DFdiscover** processes. These processes run in the background and do most of the actual work. These processes are: *master*, *slaves*, *study servers*, *EDC servers*, *inbound*, and *outbound*. Four of these processes (master, slave, EDC server and outbound) are referred to as daemons because they run continuously in the background, without human intervention, as long as the **DFdiscover** system is operational. This distinguishes them from the study database and inbound servers that are automatically started when needed and exit when their task is done.

An operational **DFdiscover** system will have all of these background processes running concurrently on the same or different computers on your network. There will only be one copy of the master and at most one copy of each database server. In order to send documents there must be an outbound daemon. And finally, if a **DFdiscover** system has recently received several documents, the inbound server may be running, dealing with one of the newly arrived documents.

Master Daemon

The master daemon is started automatically when the host server computer starts. It can also be started by executing the command

IMPORTANT: Typically, the daemons are defined at the system level so that the operating system starts them when the computer boots and halts them before the computer is shut-off. Hence, it is rare that a user will have to manually execute this command.

```
# /opt/dfdiscover/bin/DFbootstrap
```

It runs on the licensed **DFdiscover** server and has 4 primary functions:

- It controls the total number of users (or clients) that are allowed to run concurrently on the system. This is constrained by the **DFdiscover** software license.
- It serves as a **connection router** for **DFdiscover** client applications, servicing requests for a connection to a particular server. It determines where the requested server is running and returns this information to the client. If the requested server is not running, the master solicits a slave daemon to start it.
- It receives notification of each incoming document and responds by soliciting a helper application to process the incoming document.
- It manages outgoing documents (typically Query Reports being that are emailed/sent back to the clinical sites) by assigning the document to the outbound daemon.

Study Server

When a **DFdiscover** application needs information about how a study is configured (e.g. where is the default study printer located?), or needs to read from or write to a study database, it requests a connection to the study database server. The master honors this request, if it can, and if not, it solicits a slave to start the study server, making sure that there is one server process executing for each **DFdiscover** study.

It is the job of the study database server to serialize all database transactions, and to lock data records as needed, ensuring that only one user is able to modify a record at a time.

EDC server

The EDC server handles the secure, encrypted communication between an EDC client and the study server. The EDC server starts another EDC server child process automatically when a new data collection tool connection is established and stops the child process when the data collection tool client exits.

Inbound Server

The job of the inbound server is to process the PDF and TIFF documents received by email and from **DFsend**, as well as G3 or TIFF image files received by fax software. This includes:

- breaking down the incoming files into individual pages,
- de-skewing and flipping pages as necessary,
- reading the CRF barcodes to identify the study they belong to,
- running the ICR software to generate an initial (workflow level 0) data record for each CRF page,
- sending the pages and data records to the appropriate study database server, or if none is identified, to the image router,
- and finally, archiving the original document.

The inbound server is started, and assigned a document to process, by the master daemon that receives email whenever a new document has arrived. When the inbound server has finished the current document, the master will assign it another document, if there is one. Otherwise, the server exits, and will be subsequently re-started by the master when needed.

Outbound Server

The outbound server runs continuously and handles all outgoing document transmissions from **DFdiscover**. The outbound server receives requests from the master to send a document, queue the document for faxing or emailing, and track the document to its completion or ultimate failure. These documents are typically quality control reports being sent to the clinical sites participating in **DFdiscover** studies.

DFdiscover System Limits

The following is a complete listing of **DFdiscover** database limits and formats used in **DFdiscover** system administration. The same parameters may appear in other tables with different limits because the limit is different in that context.

Description	Limit	Comments
Number of outbound documents that can be queued concurrently	1-65535	
Username	16 chars maximum	alphabet A-Z, a-z, 0-9, _ (underscore)
Study number	1-255, 1-999	The suggested range for study numbers that require barcodes is 1-249. Study numbers of 250-255 may be used by DF/Net Research, Inc. in distributing example or test studies such as study 254, Acceptance Test Kit. On an appropriately licensed system, it is further possible to use study numbers in the range 256-999 for EDC studies.
Plate number	0-511	Plate 0 references the new record queue.
Workflow level	0-7	Level 0 represents new, not yet validated, or pending records. Once a record has a level greater than 0, it cannot be assigned level 0 again.

Security

DFdiscover implements product and database security at a variety of levels. To fully utilize **DFdiscover** it is important to understand the security choices available.

The implementation of a proper security model for **DFdiscover** requires understanding secure communication, authentication, what permissions are available to users, and then appropriately allowing or restricting those permissions.

This chapter presents the information necessary to implement security within **DFdiscover**. Rather than repeat the details already available in other chapters or manuals, some sections include cross-references.

Secure Communication

Within **DFdiscover** communication between client applications and server is via encrypted communication on port 443. This port must be open on any firewalls between the local computer and the study server.

The security of the communication is based upon 3 industry standard technologies:


1. Communication protocols, namely TLS 1.2 or 1.3,
2. Strong encryption ciphers, and
3. Independent certification and confirmation of the server.

This is industry-standard technology that encrypts the bi-directional communication using a 'certificate of trust' provided by the server. It is the same technology used by banks and the majority of secure, global web services.

You can visually confirm that the communication is secure by examining the details of the communication protocol and encryption cipher. In the login screen of any desktop client, click the green lock icon next to the **DFdiscover Server** input field.



You can also examine the certificate of trust. After login, select **Help** > **Certificate Info** and look for the green checkmark.



User Authentication

The first level of security in **DFdiscover** is user authentication - before a user can access **DFdiscover** they must provide a valid username and password combination. The mechanism for implementing that username and password combination depends upon the type of user and type of application:

- **DFexplore, DFweb, DFcollect, DFengage, DFsend, DFsetup and DFadmin user (client-side, interactive user):** Data collection tool users are authenticated at the **DFdiscover** application level - the username and password combination are created, maintained, and verified directly by **DFdiscover**. This has the advantage that these users do not need a UNIX login account.

Creation and administration of these accounts is performed entirely by a **DFdiscover** administrator - there is no need to involve a UNIX administrator as these users will not need access to the underlying UNIX system.

- **DFexport, DFbatch, DFattach and DFpdfpkg user (client-side, command-line user):** Like the interactive user applications, these applications are also authenticated at the **DFdiscover** application level. As a preliminary step however, the user must use their operating system credentials to gain access to a command-line.

Since there is no login dialog for these applications, the user is required to supply their authentication information through command-line options or environment variables. Users can also choose to manage their password information locally on their desktop computer with **DFpass** see [Programmer Guide, DFpass](#) and [Programmer Guide, User Credentials](#). Proper implementation of **DFpass** allows users to access command-line applications in scripts and the cron facility without exposing their password as clear text.

- **Central data management office DFdiscover user (server-side user):** This is the "traditional" **DFdiscover** user that is authenticated by a username and password combination created, maintained, and verified at the UNIX operating system level.

Each such user is identified by a unique username. That **DFdiscover** username is also the UNIX login name. Each UNIX login name also has a password, and that combination of login name and password uniquely identifies an individual within the operating system, and also within **DFdiscover**.

Before this type of user can access **DFdiscover**, they require a login account which is defined by the UNIX administrator. In the process of creating this login account, a unique login name and password are assigned. The password is the responsibility of the user and should be changed regularly, in accordance with the company's computer security requirements. A user's login name is publicly available so it is critical that each user maintain the privacy of their password.

Once a user has a UNIX login name and password, they can login to the UNIX system. Once they are logged in, they also have general access to **DFdiscover** and are identified in **DFdiscover** with their UNIX login name, also known as their **DFdiscover** username. However, each user is required to have a **DFdiscover** login account so permissions can be granted to each user for accessing various **DFdiscover** applications and studies. The **DFdiscover** login account is administered by the **DFdiscover** administrator (or) study administrator using **DFadmin**.

Two-Factor Authentication

The **DFdiscover** administrator may have enabled two-factor authentication for some or all login accounts. If two-factor authentication is enabled for a user, after successful authentication the login screen updates to request a security code (see [DFexplore User Guide, Two-Factor Authentication](#)). Independently, the user will receive an email or SMS text message, depending on system and user configuration (see [Azure SMS Messaging](#) and [User Profile](#)). The email or text message contains a unique 6-digit security code. This security code must be entered in the login dialog to complete login verification. Each security code is unique to a login username on a specific device and expires if not used within 10 minutes.

Each verification in two-factor authentication, with the same credentials and device, is valid for 30 days. After 30 days the process will repeat and a new security code will again be needed. Logging in with different credentials and/or a different device will also initiate two-factor authentication. Deleting or deactivating a user account will reset the validity of all of the user's previously used devices. Upon reactivating any user account a new security code must be entered in the login dialog to complete login verification on each device.

Microsoft Entra ID Single Sign-On (Formerly Azure AD)

NOTE: Microsoft Entra ID Single Sign-on capabilities are only available when connected to a **DFdiscover** server and API version 5.7 or later.

Microsoft Entra ID allows users to log in to **DFweb** services using their Microsoft credentials. Users can opt to utilize this method as an alternative to the default **DFdiscover** login methods which simplifies the authentication process and enhances the security of **DFdiscover**. Authentication requirements are set within the Microsoft Active Directory environment and enforced by the **DFedcservice**. The use of Microsoft Entra ID does not block the ability of users to login via **DFdiscover** using the current login and password credentials but rather serves as an alternative login method. Further information about Microsoft application registration documentation can be found here: <https://learn.microsoft.com/en-us/azure/active-directory/develop/quickstart-register-app>

Configuring the use of Microsoft Entra ID services with **DFedcservice** requires the registration of an application instance within Azure Active Directory. This establishes a trust relationship between **DFdiscover** API services and the identify provider, the Microsoft identity platform. During the **DFdiscover** user account creation process a valid email address is required to utilize the Single Sign-On feature.

Configuring the **DFdiscover** API services for Single Sign-On requires the following settings during the application registration process:

- **Single-Page Application** Redirect URI. The URIs we will accept as destinations when returning authentication responses (tokens) after successfully authenticating or signing out users. The redirect URI you send in the request to the login server should match one listed here. Also referred to as reply URLs. Replace the `dfweb.dfdiscover.com` FQDN with your local **DFweb** instance FQDN.

```
https://dfweb.dfdiscover.com/login/dfwebssso
https://dfweb.dfdiscover.com/common
```

- **Select the tokens you would like to be issued by the authorization endpoint:**

- Check (enable) **Access tokens (used for implicit flows)**
- Check (enable) **ID tokens (used for implicit and hybrid flows)**

To enable Single Sign-On for DFedcservice the `/opt/dfdcover/lib/DFedcservice.cf` file. The DFedcservice file requires both the use of `azureclientid` and the `azuretenantid` which is a crucial setting in the configuration file and can take one of the following values: `organizations`, `common`, or a specific `tenantID`.

- **organizations:** This allows users from all organizations associated with the Microsoft Entra ID (Azure AD) tenant to access the application.

```
master apache.dfnetwork.com
key /opt/dfdcover/lib/edckey.pem
certificate /opt/dfdcover/lib/edccert.pem
bindaddr 192.168.3.30
ciphers ECDHE-RSA-AES128-SHA,ECDHE-RSA-AES128-SHA256,TLS_AES_128_GCM_SHA256,TLS_CHACHA20_POLY1305_SHA256
azuretenantid organizations
azureclientid 11111111-1111-1111-1111-111111111111
```

- **common:** This allows users with both personal Microsoft accounts (MSA) and work or school accounts from any organization associated with the Microsoft Entra ID (Azure AD) tenant to access the application.

```
master apache.dfnetwork.com
key /opt/dfdcover/lib/edckey.pem
certificate /opt/dfdcover/lib/edccert.pem
bindaddr 192.168.3.30
ciphers ECDHE-RSA-AES128-SHA,ECDHE-RSA-AES128-SHA256,TLS_AES_128_GCM_SHA256,TLS_CHACHA20_POLY1305_SHA256
azuretenantid common
azureclientid 11111111-1111-1111-1111-111111111111
```

- **tenantID:** Replace this with the actual tenant ID of the specific organization you want to allow access. This setting restricts access only to users from that particular organization.

```
master apache.dfnetwork.com
key /opt/dfdcover/lib/edckey.pem
certificate /opt/dfdcover/lib/edccert.pem
bindaddr 192.168.3.30
ciphers ECDHE-RSA-AES128-SHA,ECDHE-RSA-AES128-SHA256,TLS_AES_128_GCM_SHA256,TLS_CHACHA20_POLY1305_SHA256
azuretenantid 00000000-0000-0000-0000-000000000000
azureclientid 11111111-1111-1111-1111-111111111111
```

NOTE: During user account creation in **DFadmin** uniquely defined email addresses can be associated with only one user account. Duplicate user emails cannot be assigned across various login accounts.

DFedcservice logs all SSO sign-ons in the `/var/log/messages` file

```
Jul 28 10:24:43 explore dfweb: #033[0mGET /blank.html #033[36m304#033[0m 0.415 ms - #033[0m
Jul 28 10:24:44 explore dfedcservice[17922]: new connection from explore
Jul 28 10:24:44 explore dfedcservice[17922]: https://login.microsoftonline.com/organizations/v2.0/.well-known/openid-configuration -> 1589 bytes
Jul 28 10:24:44 explore dfedcservice[17922]: https://login.microsoftonline.com/organizations/discovery/v2.0/keys -> 7780 bytes
Jul 28 10:24:44 explore dfedcservice[17922]: testuser@explore /DFedc?
cmd=authorize&username=&password=&encoded&prog=DFws&apiprogram=DFweb&protocol=5.7.0&json
heartbeat=3&device=74d15730-749c-11ed-949f-e570dee3d1f1 -> 561 bytes
Jul 28 10:24:44 explore dfedcservice[17922]: user testuser@explore logged in using DFweb with Microsoft SSO email testuser@dfnetwork.com.
```

Restrict Login to Microsoft Entra ID SSO Only

For organizations that wish to only allow users to log in with Microsoft Entra ID SSO, this can be enabled per server in the DFedcservice configuration file. This setting will remove the username and password fields from login screens to prevent logins with **DFdiscover** credentials, and only allow login with Microsoft SSO. This setting is enforced in all desktop applications as well as **DFweb** and **DFcollect**.

To enable SSO-only login, add the `ssologinonly` parameter with a value of 1 to the `/opt/dfdcover/lib/DFedcservice.cf` file. If you wish to later disable SSO-only login, change the `ssologinonly` parameter value to 0 to allow logins with **DFdiscover** credentials as well as with Microsoft Entra ID.

Azure SMS Messaging

NOTE: Azure SMS Messaging is **disabled by default** and **region-limited** at the server level. It is only available when enabled by the system administrator and when connected to DFdiscover server and API version 5.11 or later.

DFdiscover provides a region-limited option to send text messages via SMS to users using Azure Communication Services, when enabled for the server. SMS messaging is configured at the server level, with all studies and users on that server using the same, region-specific Azure SMS service and sender.

When enabled for a server, SMS is available as an option for two-factor authentication (2FA), delivering authentication codes by text message instead of email based on the user-specific setting defined in the **DFadmin User Profile** by a DFdiscover or study administrator. Phone numbers are managed in the user profile and may be updated by the user in **DFexplore** or **DFweb**. For ePRO users, phone numbers are managed by ePRO Admins in **DFweb** (see [ePRO Admin Permissions](#)). System and study administrators should ensure appropriate policies and controls are in place when enabling SMS-based authentication.

Each DFdiscover server is configured to send SMS messages using a single Azure SMS service and sender configuration. As a result, SMS delivery is limited to phone numbers within the regions supported by that configuration. If SMS 2FA is enabled for users with phone numbers outside the supported region, those users will not receive authentication codes and will be unable to log in.

When SMS is enabled for a server, by default DFdiscover supports SMS delivery to phone numbers in North America (United States, Puerto Rico, and Canada) only. Individual servers may be configured differently, including by organizations hosting their own DFdiscover environment and using their own Azure SMS service.

SMS configuration is enabled and defined in the DFedcservice configuration file, including the Azure service connection and message template. For more details, see [DFdiscover System Files - DFedcservice.cf](#).

NOTE: If SMS is disabled in DFedcservice.cf when the 2FA **Send code by Azure SMS** option remains enabled for any users, those users will receive their 2FA codes by email. If the email address is not defined for those users, the user will see the error message "No email address found for two-factor authentication" and will not be able to log in.

Server Filesystem Permissions

The UNIX operating system restricts access to each file and directory by ownership:

- **owner**: Is the user the owner of the file? The person that creates any new file or directory becomes the owner of it. Ownership cannot be changed except by a UNIX administrator.
- **group**: Is the user in the same group as the file's group? Each user belongs to one or more groups, their primary group and zero or more secondary groups. Groups exist to simplify sharing of files among collaborating users and are created and assigned by the UNIX administrator. The default group for **DFdiscover** is studies. Unless a finer segregation of groups is required for **DFdiscover**, all **DFdiscover** users should have studies as their primary group.

When a file or directory is created by a user it is assigned that user's primary group for group ownership.

- **other**: If the user is not the owner of the file, and they are not in the same group as the file's group, they are considered to be part of the general population of users, with no special ownership privilege for this file or directory.

and action:

- **read**: Is permission granted to read the file?
- **write**: Is permission granted to write (overwrite, update, or append) to the file?
- **execute or search**: Is permission granted to execute the file (meaningful only if the file is an executable application), or to search inside the directory?

To ensure that files which need to be shared among **DFdiscover** users can be easily shared it is important to understand and properly implement filesystem permissions. To assist in the consistent application of filesystem permissions, the UNIX-provided **umask** setting should be used. The **umask** setting can be set or altered at any time but it is most advantageous when it is specified in the user's login file, their .profile or .login, and then applies consistently while they are logged in. The default **umask** setting is 022, which prevents a file created by a user from being updated by any other. This setting can hinder collaboration within **DFdiscover** and it is recommended that the default setting for **DFdiscover** users be 002.

There is additional information about filesystem permissions at the study database level in [Study Filesystem Permissions](#), at the end of this chapter.

A full description of UNIX filesystem permissions is beyond the scope of this document but is easily found in any UNIX administration guide.

PDF Encryption Password

At the user-level, it is possible to encrypt, using standard PDF encryption, any or all PDFs created by the user. This is not a permissions related issue per se, but is relevant to user login settings and hence to this section.

PDFs can be optionally encrypted, encryption providing an additional security layer if the PDFs are transmitted electronically. Encryption occurs at PDF creation time using a one-time encryption password provided by the user, or using the user-specific encryption password found in the file .dfpdfpasswd in the PDF creator's home directory. This file must contain exactly one plain-text line, and this line is read and used as the user's encryption password during PDF creation.

IMPORTANT: File permissions on this file should be set to 0600, preventing other users from viewing or modifying the file contents.

DFdiscover System-Level Permissions

Any user that can successfully authenticate on a UNIX host which is running **DFdiscover**, and has access to a command-line, can:

- access individual study databases to which they have been granted permission,
- access the data collection tool client application running on the server,

- access the **DFsetup** client application running on the server,
- view **DFdiscover** license activity
- view the **DFdiscover** documentation.

Access to individual studies is controlled by the **DFdiscover** and/or study administrator and is further detailed in [Users](#) and [Roles](#). It is not possible for any user to access any study database unless they have been explicitly granted permission by an administrator. Permissions to access data collection tools and **DFsetup** can be altered for each user as indicated in [Permissions](#).

DFdiscover Study-Level Permissions

Study roles, as defined and assigned to users in **DFadmin**, determine a user's permission to use the **DFdiscover** client applications and determine what they are able to do within each of these applications. The permission specifications provided by **DFdiscover** can be very detailed and are described in [Roles](#).

Impact of Database Access Restrictions on Standard Reports

Database access restrictions can influence the output of each report. If database access restrictions have been implemented by the **DFdiscover** administrator for the study, execution of the same report by different users may lead to different results, and some users may not be able to run certain reports at all. The table summarizes the effect that access restrictions have on each report.

Conformance with Access Restrictions for Standard Reports

Conformance	Report Name(s)
(none) : All database access restrictions are ignored. The output from this report is not affected by access restrictions and is the same for all users.	DF_ATcrfs, DF_ATfaxes, DF_ICcenters, DF_ICschema, DF_ICvisitdates, DF_ICvisitmap, DF_QCfaxlog, DF_QCstatus, DF_SScenters, DF_SSschema, DF_SSvars, DF_SSvisitmap, DF_WFcrfs, DF_WFcrfsperwk, DF_WFdiskusage, DF_WFfaxes, DF_WFqcs
(record) : All database access restrictions are applied. The output is restricted to those records for which the user has access permissions. The output from this report will vary across users and will reflect their individual record access restrictions.	DF_ATmods, DF_ICkeys, DF_ICqcs, DF_ICrecords, DF_PTCrfs, DF_PTlist, DF_PTmissing, DF_PTqcs, DF_PTunexpected, DF_qcsbyfield, DF_stats
(subject) : All database access restrictions are applied. The output is restricted to those subjects for whom the user has full visit and plate access permissions. The output from this report will vary across users and will reflect their individual subject access restrictions.	DF_PTvisits
(site) : All database access restrictions are applied. The output is restricted to those sites for which the user has full subject, visit and plate access permissions. The output from this report will vary across users and will reflect their individual site access restrictions.	DF_CTCrfs, DF_CTqcs, DF_CTvisits, DF_PTschedule, DF_QCfax, DF_QCprint, DF_QCreports, DF_QCsent, DF_QCview
(full) : All database access restrictions are applied. Only users with unrestricted access to all records in the study database can run this report. The output from this report will be the same for all such unrestricted users.	DF_ICimages, DF_QCupdate, DF_XXkeys, DF_XXtime

Impact of Database Access Restrictions on Edit Checks

In the data collection tools the data fields available to edit checks are restricted by the user's site, subject, visit, plate and read level access permissions. This helps to ensure that users will not be shown database values in edit checks to which they would normally not have access, but it also means that data records unavailable to the user cannot be used in edit checks. This restriction does not apply to hidden fields. Thus data values that are needed in an edit check, but should not normally be seen by users with a certain role, can be made available by defining them as hidden fields on plates to which the user has read access permissions.

In **DFbatch** edit checks have *unrestricted access to all study data* and ignore all access restrictions on the current user. To mitigate any negative consequences of allowing edit checks unrestricted access to all study data in **DFbatch**:

- edit check programmers must be aware of database access restrictions for the study and design their edit checks to respect the intentions of the study management team. Although an edit check may require access to restricted data to perform some logic check, the programmer should be careful not to display the restricted information to users not allowed to see it.
- permissions on the study edit checks file should be maintained as restrictive as possible so that only authorized edit check programmers can access the file and add or update the edit checks.

ePRO Scheduling Permissions

A special system user *dfeproagent* is required for studies using **DFengage** ePRO and providing access to the **ePRO Management** page in **DFweb** via [ePRO Admin Permissions](#). The *dfeproagent* user is used to determine the subject ePRO statuses in **DFweb** and identify due and overdue activities for reminders sent by [DFeproreminders](#). The user account requires the following:

- Access to all sites and subjects in each ePRO study with a role that includes access to **DFweb** and provides the same database permissions used by ePRO participants to ensure consistent visit scheduling. Apply the same ePROdata role used by participants to the *dfeproagent*, or create a separate role.
- Password set and not expired.

- Password registered with **DFpass** (see [Programmer Guide, DFpass](#)).

Study Filesystem Permissions

Many tasks in **DFdiscover** rely upon successful access to one or more files in the underlying filesystem. Given that the study is administered in the manner described in this chapter, no special attention needs to be paid to the permissions on those underlying files. However, it is still recommended that the **DFdiscover** administrator periodically review and update the permissions on the filesystems accessed by each **DFdiscover** study using [Programmer Guide, DFstudyPerms](#). This practice is further detailed in [Monitoring study directory permissions](#).

Getting Started with DFadmin

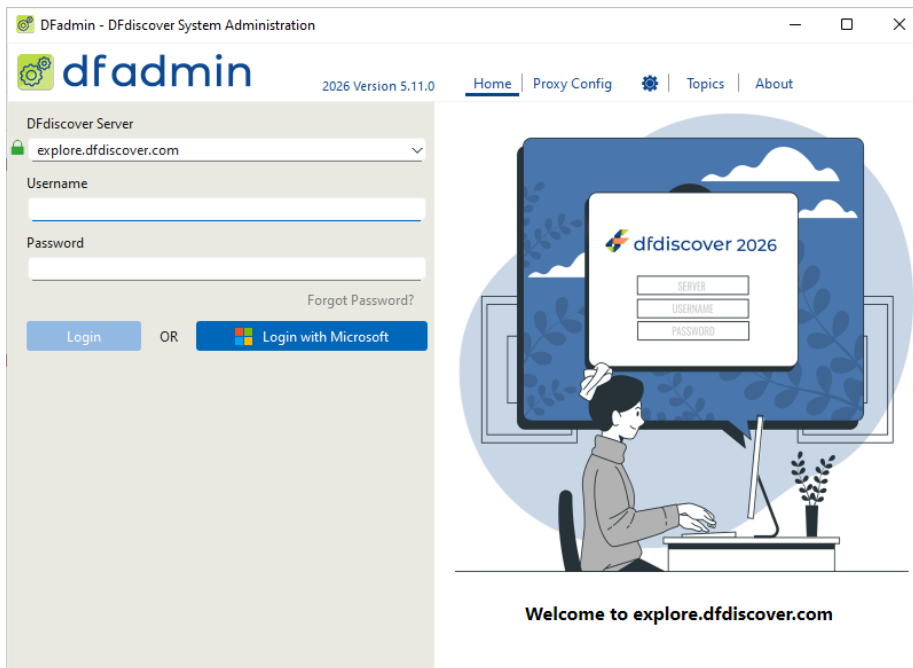
Starting DFadmin

To start **DFadmin**:

- On **Windows** and **macOS**, click (or double-click) the **DFadmin** icon.
- On **Linux**, type the following at a shell prompt:

```
# /opt/dfdiscover/bin/DFadmin
```

Once the application is launched, the login screen is presented.



The **DFdiscover Server** field requires the hostname of the **DFdiscover** server. You must enter your valid login credentials in the **Username** and **Password** fields. In addition, you must have been previously granted **DFdiscover** or study administration permissions.

NOTE: Some organizations install a proxy server to filter requests sent to the internet. If your computer is behind a proxy server click **Proxy Server** to configure the proxy server. Contact your IT department for the necessary information.

If Single Sign On with Microsoft is enabled on your **DFdiscover** server, a [Login with Microsoft](#) button is shown. You can log in with your Microsoft account instead of using your **DFdiscover** username and password. Click [Login with Microsoft](#) to begin the Microsoft login process.

In some cases, [Login with Microsoft](#) may be the only method of login available, depending on how the administrator has configured the **DFdiscover** server.

NOTE: If you log in using a Microsoft login, any password requests during your **DFadmin** session will require you to authenticate using the same Microsoft login process.

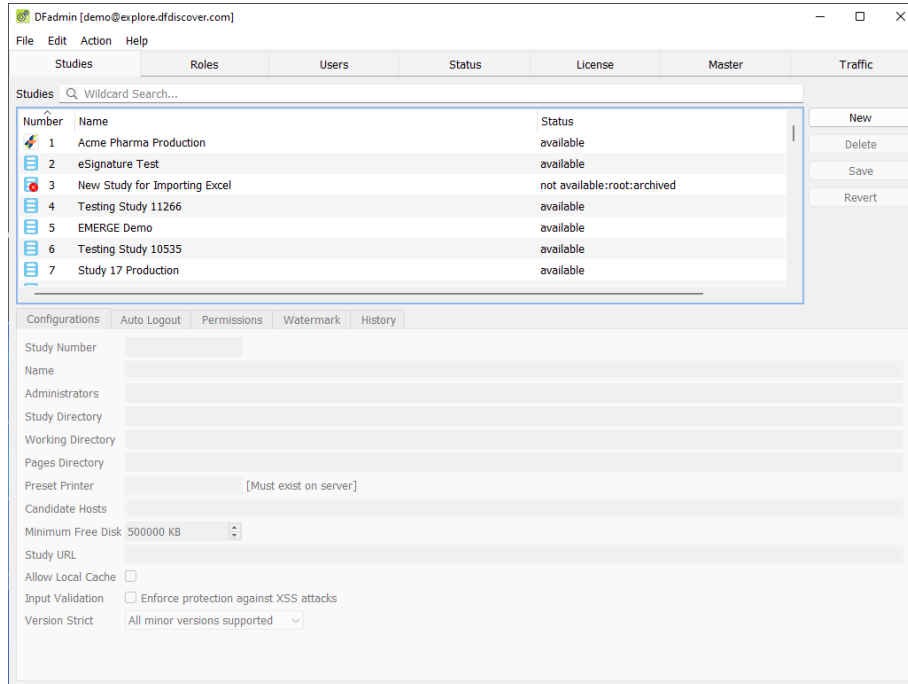
If you forget your password, you can request a password reset from the login dialog. First enter the server name and your username into the login dialog, then click **Forgot Password?**. You are asked for your email address, which must match the email address for you on this server. A link to reset your password is emailed to you. Click the link and follow the instructions to set a new password, then return to the login dialog to log in with your new password. This feature is only available if it has been allowed by the **DFdiscover** administrator.

After successful authentication, the **DFadmin** application appears.

The user guide is available by selecting [Help](#) > [Topics](#) from the application menu after logging in. Once selected, the user guide will open in the default web browser.

DFdiscover System Administration Tasks

DFadmin provides tools for DFdiscover and study administrators to perform common system administration tasks using a tabbed visual interface.



The views available in DFadmin are:

Studies

- register new studies, update existing studies
- identify location of the study directory and files relevant to the study configuration
- review permissions of study users

Roles

- define permissions and structure into meaningful roles
- set maximum workflow level and application permissions for each study

Users

- create user accounts for staff that are allowed to use DFdiscover
- administer all activities relevant to user account management including contact information, password and account status
- assign study roles to users

Status

- check DFdiscover services status
- enable/disable studies
- review study connection details

License

- manage DFdiscover password needed to activate software license
- review license usage statistics

Master

- register email address to receive automated warning/error messages
- specify login password complexity, reuse and notification rules
- identify the default printer
- specify a text banner to appear on login dialogs

Traffic

- define server available to process incoming documents
- define server available to process outgoing documents
- configure **DFdiscover** to send back a receipt when documents arrive from selected clinical sites

License

Introduction

DFdiscover is licensed by a unique host identifier for a single server machine. This licensed machine is the only machine permitted to execute the **DFdiscover** master daemon. This machine is often referred to as the "**DFdiscover** master".

DFdiscover uses a concurrent user license model, which permits up to a purchased maximum number of concurrent users to be using the software, and does not limit which users are running the client applications, or where those applications can be run. Only the **DFdiscover** master daemon needs to be licensed; the individual client machines do not.

If you are licensing **DFdiscover** for the first time, you must choose your **DFdiscover** master carefully. In particular, this machine should have demonstrated network and hardware reliability. If the **DFdiscover** master is 'down' (unavailable), **DFdiscover** will not function on that machine or any client machines.

DFadmin will only connect to a server that already has a valid license. To enter license information for new installations or to update license information on servers where the license has expired, use **DFserveradmin** (see [Software Installation Guide, Request and Install DFdiscover License](#)).

Each **DFdiscover** license is controlled by a password which encodes 4 attributes:

- the unique host identifier of the master,
- the license expiration date,
- the number of license seats, and
- optional features.

These 4 attributes are provided to DF/Net Research, Inc.. Using these attributes as input, a password is computed and provided back to you. Since the provided password encodes the 4 attributes, changing the value of any attribute without also requesting and updating the password invalidates the software licensing, potentially making **DFdiscover** unusable.

DFadmin's License Tab

DFadmin provides a view for updating license information, as well as a view for reviewing current license usage.

The **License** tab has two views: **License** for entering **DFdiscover** software license details, and **Usage Statistics** for reviewing seat usage.

Entering and Reviewing License Information

The **License** view is where all changes to license information are made.


Click the **License** tab. The **License** view is separated into two sections, separated by a horizontal rule: an upper section for values that are password restricted, and a lower section for values that can be modified independent of the password.

Password Restricted Values


The following values appear in the upper section:

- **Host Identifier:** The host identifier is a 20 character alphanumeric value expressed in 5 dash-separated blocks of 4 characters/digits. It is unique to one computer. For **DFdiscover**, this is the host identifier of the server machine that has been chosen to run the master daemon. This value does not change and should not be modified.
- **Total Seats:** This number defines the maximum number of concurrent user logins that are permitted. The number, in the range 1 to 500 inclusive, is equal to the number of seats purchased. Normal use of the software will occupy seats up to the maximum number of concurrent logins permitted.
- **Expiration Date:** **DFdiscover** is licensed for a defined period of time. The software license expires at midnight on the license expiration date.
- **Features:** **DFdiscover** optionally includes features that may be valuable to a certain type of client. These features can be licensed separately. Features are defined by an acronym and an optional seats limiter. Features available for licensing are **DFWS** (web services via API), **LS** (large study numbers), **DFENGAGE** (ePRO mobile app), and **DFCOLLECT** (online/offline EDC mobile app).

Individual features can be directly entered in the field, or they can be selected from the dialog. Multiple features are separated by a comma (,). The maximum number of seats available to and **DFWS** connections can be limited, whereas the other features (**LS**, **DFENGAGE**, **DFCOLLECT**) are simply available or not.



The screenshot shows a dialog box titled "Set feature's seats" with a question mark icon and a close button. It contains a dropdown menu for "Feature" set to "DFWS", a text input field for "Seats" containing the number "40", and an unchecked checkbox for "Maximum". At the bottom are "OK" and "Cancel" buttons.



The screenshot shows a dialog box titled "Set feature's seats" with a question mark icon and a close button. It contains a dropdown menu for "Feature" set to "LS", an empty text input field for "Seats", and a checked checkbox for "Maximum". At the bottom are "OK" and "Cancel" buttons.

- **Password:** The value for this field is provided by DF/Net Research, Inc. on payment of the software license fee. The value is 20 characters in length, presented as 5 dash-separated blocks of 4 characters or digits. To prevent confusion, the digits 0 ('zero') and 1 ('one'), and the letters O ('oh') and I ('eye') never appear in a password.

Password Independent Values

The following values appear in the lower section, below the horizontal rule. They can be updated independent of the password value.

- **Hostname:** The hostname is the name by which the **DFdiscover** server is known on your network, and is the output of the command:

```
% hostname
```

This field is not used when defining/evaluating the password, however it is verified in other areas of the software. Be careful before changing it.

- **Admin Seats:** Seats can be reserved for use by **DFdiscover** and study administrators, to ensure that they are never locked out because all licensed seats are taken by users. We recommend that at least 1 seat be reserved for administrators. The split between administrative seats and user seats can be changed at any time. Each login by a **DFdiscover** or study administrator takes an administrative seat if one is available; if an admin seat is not available, administrator login will try to take a general user seat. All other users cannot take one of the admin seats even if one is available and all user seats are in use.

It is also possible to change the allocation of seats assigned to a feature (as a subset of the purchased **Total Seats**). Specifically, it is possible to assign an upper limit to seats that can be used by **DFWS**. The notation is **FEATURENAME:MAXSEATS**. For example, **DFWS:40** indicates that feature **DFWS** has been licensed and that a maximum of 40 seats can be occupied by users/clients using that feature.

This change can also be made independent of a password change.

DFadmin [demo@explore.dfdiscover.com]

File Edit Help

Studies Roles Users Status License Master Traffic

License Usage Statistics

Host Identifier QV7H-B9AD-L3CM-ABL7-GZ48

Total Seats

Expiration Date

Features

Password

Hostname explore.dfdiscover.com

Admin Seats User Seats 500

Save Revert

Applying Changes

Verify that each field has been entered correctly. To save and apply the license information, click **Save** or select **File** > **Save**. If any license fields are incomplete or incorrect, the save will fail and the existing license will be left unchanged.

If complete and accurate, the license information is saved and passed to the **DFdiscover** master daemon. It is not necessary to restart **DFdiscover** - the expiration date, seat count and features on a currently licensed system will be automatically updated.

Subsequently, each time the **DFdiscover** master starts, it verifies the license. The master (and hence **DFdiscover**) will fail to start if the license file cannot be found or any of the information in the file is incorrect.

Click **Revert** to undo changes you have made to any fields in the license view.

License Expiry

DFdiscover warns of approaching license expiry [60, 30, 14, 7, 6, 5, 4, 3, 2, and 1] days prior to the expiry date. The person recorded as the **Problem Mail Recipient** (in the Master tab) will receive email messages similar to the following:

```
date machine DFmaster.rpcd[pid]: your software license expires in # days
```

If the license is not renewed, it will expire at midnight on the expiry date. In this case the essential **DFdiscover** services will halt and not restart, and the **Problem Mail Recipient** will receive an email message similar to the following:

```
date machine DFmaster.rpcd[pid]: your software license has expired
```

If the **DFdiscover** license expires, **HylaFAX** will continue to receive CRF images. When the **DFdiscover** license is subsequently renewed, the received CRF images will be processed by **DFdiscover**. Nothing will be lost. Be careful however that the incoming images directory has sufficient disk capacity to store any CRF images that arrive while the **DFdiscover** software is off-line.

Updating License Information without DFadmin or DFserveradmin

It is possible to apply new license information to an active **DFdiscover** installation without the use of **DFadmin** or **DFserveradmin**. Logged-in to the licensed **DFdiscover** server, the system super-user can update/install a new license file and send a **UNIX** SIGHUP signal to **DFmaster.rpcd**. The SIGHUP signal is interpreted as a request to re-read the license file. Before sending the SIGHUP signal, the new license information must be entered into the **DFdiscover** license file, namely `/opt/dfdiscover/work/license`.

From the command-line, the following steps cause **DFmaster.rpcd** to re-read its license:

1. Before editing `/opt/dfdiscover/work/license`, make a backup copy of the file.
2. Edit `/opt/dfdiscover/work/license` with a text editor and save the changes.
3. Determine the process id, *pid*, of **DFmaster.rpcd**:

```
# ps ax | grep DFmaster
```

4. Send the SIGHUP signal to **DFmaster.rpcd**:

```
kill -HUP pid
```

If the new license can be loaded and validated successfully, the new settings are applied. Otherwise, they are ignored. The interval to license expiry is re-calculated if necessary which may also reset the timing for the license expiry warning emails.

WARNING: Invalid License Settings If the new license file cannot be loaded and validated successfully, it is imperative that the old license settings be re-entered and saved to the license file. Although invalid settings are ignored in this scenario, valid settings are required if there is a need to restart **DFdiscover**.

A recommended practice is to keep a backup copy of a valid license file before attempting to update the active license file from the command-line.

Reviewing License Usage Statistics

DFdiscover provides current and cumulative license usage information in the **Usage Statistics** view.

	Limit	In Use	Max Used	Rejected
Total Seats	500	2	7	0
User Seats	500	2	7	
Admin Seats	0	0	0	
Features				
DFCOLLECT	500	0	5	0
DFENGAGE	500	0	0	0
DFWS	500	1	2	0
LS				

Statistics are available for seat usage and feature usage. Seat usage is reported by **User Seats** and **Admin Seats**, mirroring the specification in the license itself. Feature usage is reported individually for each licensed feature.

For seats and features, **DFdiscover** is able to report the licensed value (**Limit**), the current value (**In Use**), the maximum value since start (**Max Used**) and the number of license requests that could not be fulfilled (**Rejected**).

The statistics reported include:

Statistics since date/time	The date and time of the beginning of the reporting period. Reported usage is available since DFdiscover last started, or the license was last updated.
Limit	The number of seats included in your current DFdiscover license is shown in the first row of the table. This is a static value matching the value of Total Seats in the License view.
In Use	This is the number of concurrent seats currently in use by clients connected to the DFdiscover server. The number of seats available is the difference between Limit and In Use .
Max Used	The maximum number of seats that have ever been used concurrently is reported in this column. If the maximum has been reached (matches the limit) and you see many rejections in the next field it may be time to purchase additional seats.
Rejected	The number of times that a license request could not be granted because all available licenses were in use is reported in this column. DFdiscover does not track if rejected requests came from administrators or users.

The statistics can be updated at any time by clicking [Refresh](#).

Master

The master daemon is the most important application in the **DFdiscover** system. It connects every client application to the appropriate study database server, and without it no useful work can be done. The master is started when **DFbootstrap** is executed (which typically occurs when the licensed computer is started) and runs continuously.

NOTE: The term daemon indicates that the program runs continuously in the background. It is not under direct user control.

The master has 4 functions:

- It controls the total number of users (or clients) that are allowed to run concurrently on the system. This is constrained by the **DFdiscover** software license.
- It serves as a *connection router* for **DFdiscover** client applications, servicing requests for a connection to a particular server. It determines where the requested server is running and returns this information to the client. If the requested server is not running, the master solicits a slave daemon to start it.

- It receives notification of each incoming document and responds by soliciting a helper application to process the incoming document.
- It manages outgoing documents (typically Query Reports being that are emailed/sent back to the clinical sites) by assigning the document to the outbound daemon.

Starting DFdiscover

In most installations, **DFdiscover** and the master are configured to start automatically whenever the computer (re-)starts and stop automatically whenever the computer halts. It is also possible to start and stop **DFdiscover** manually or with the **DFserveradmin** application (see [DFserveradmin](#)). **DFdiscover** can be started manually, by *root*, with the **DFbootstrap** command as in:

```
# /opt/dfdiscover/bin/DFbootstrap
DFmaster.rpcd: your software license expires in 19 days
1. Checking state of incoming daemons...

2. Checking /opt/dfdiscover/work/.dfincoming_work...

3. Checking /opt/dfdiscover/incoming...
  0 faxes awaiting processing.
DFmaster.rpcd: started.
Done.
```

There is no further direct user interaction with the master process. It operates as a background process, receiving and directing requests from other **DFdiscover** processes.

Stopping DFdiscover

DFdiscover runs continuously on the server. It is started when the host computer is started and is halted when the computer is halted.

Occasionally it may be necessary to halt **DFdiscover**, to perform a system administration task, like installing a system upgrade, updating the **DFdiscover** software, or to install additional disk storage. Stopping **DFdiscover** can be performed with the **DFserveradmin** application (see [DFserveradmin](#)), or manually from the command-line.

DFdiscover can be halted, by *root*, with the **DFshutdown** command as in:

```
# /opt/dfdiscover/bin/DFshutdown
```

Any attempt to halt **DFdiscover** will fail if there are open client connections. In such a case, messages similar to these will be displayed.

```
Locating DFdiscover slave daemons...
Trying to shutdown slave daemons...
Slave daemon on explore.dfdiscover.com shutdown.
Done.

Trying to shutdown study servers...
Server 007 on explore.dfdiscover.com will not exit - Study server is in use.
Failed.

DFdiscover shutdown failed. Some processes still running.
```

If this occurs, it will be necessary to have all users exit their **DFdiscover** applications before shutdown can proceed.

Forcing DFdiscover to shutdown

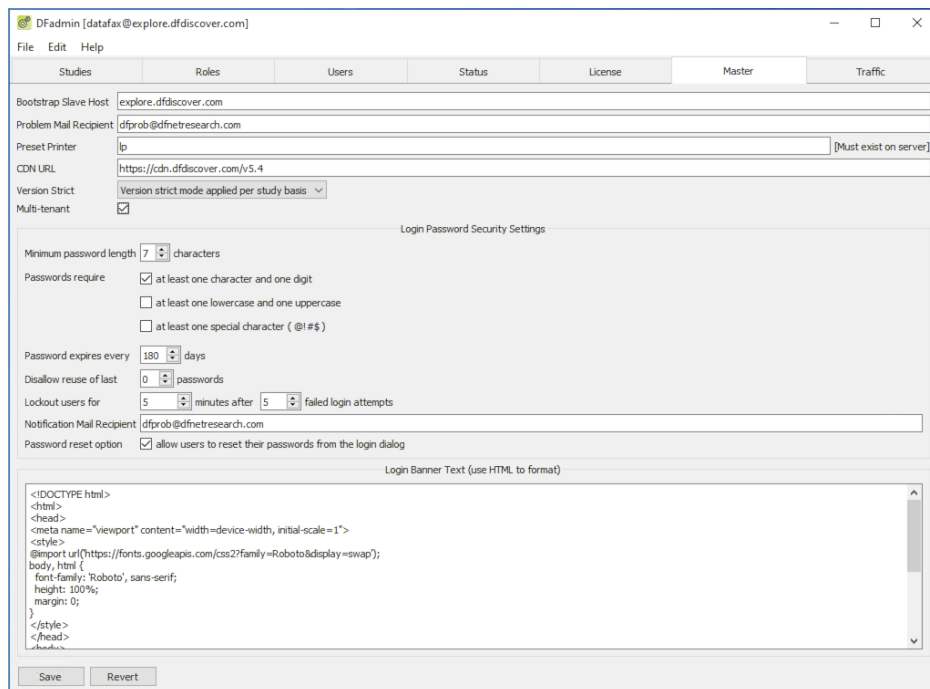
Shutdown can be forced to proceed, regardless of any warnings, by executing **DFshutdown** and including the *-f* option. This is not recommended as a general solution but instead should be used in those cases where shutdown must proceed, for example, because the computer is also halting. The invocation is:

```
# /opt/dfdiscover/bin/DFshutdown -f
```

Configuring the DFdiscover master

The master is configured with defaults during the initial installation of **DFdiscover**. In most environments, these defaults require no further modification or specification. If necessary, re-configuration of the master is simple and in most situations has to be done only once. Configuration is done through the master tab in **DFadmin**.

To (re-)configure the master, click the **Master** tab in **DFadmin**. The **Master** view has the following appearance.



The following configuration parameters can be specified in the dialog:

- **Bootstrap Slave Host:** The value for this field is typically the same as the hostname of the master, but it can also be a comma-delimited list of hostnames on the local network. When the master daemon is started, it starts one slave process on each listed host. Since **DFdiscover** server processes will only execute on machines running a slave process, each machine that will be a database server must also be listed as a slave. Good choices for database servers are machines that have local disk space (so they can swap efficiently), sufficient CPU power, and typically have a light processing load.

DFdiscover attempts to distribute work evenly to all slave machines. Also, if one machine is down **DFdiscover** will continue and try to start the server on another machine.

The default value for this field is the hostname of the computer running the master daemon. In most environments, this is sufficient.

NOTE: When attempting to distribute the load, **DFdiscover** only considers the number of **DFdiscover** servers that are running on each candidate computer. It does not consider other software that might be running.

- **Problem Mail Recipient:** Messages related to license expiry or usage will be emailed to the specified address(es). Specifically:
 - When the number of days to license expiry is: 60, 30, 21, 14, 7, 6, 5, 4, 3, 2 and 1, a warning message will be emailed
 - When license expiry occurs, an error message will be emailed
 - When no licenses are available, and a user requests a license, a warning message will be emailed
- **Preset Printer:** Several applications permit printing. Applications that are connected to a study server use the study specific value for a printer. Other applications, such as **DFadmin** or Image Router do not have a study connection. They use this value as their default print location, unless otherwise overridden.
- **CDN URL:** **DFdiscover** reports are built from JavaScript (JS) and Cascading Style Sheet (CSS) resources that reside on a separate server. By default, **DFdiscover** provides and shares a common server to make these resources available. This server is known as the CDN and the resources are rooted at the CDN URL. One advantage of this structure is that **DFdiscover** reports which are shared externally can also access these same resources via the CDN URL.
- **Version Strict:** It is possible to enforce software version numbers between client and server applications at the master level. This enforcement determines the behavior at client connection time.

NOTE: Due to important security improvements, all non-API client connections which do not support TLS v1.3 will be unable to connect to **DFdiscover** server version 5.8 and later. This includes **DFdiscover** 5.1 desktop tools. Even when the Version Strict setting allows for 5.1 clients, **DFdiscover** server 5.8 and later will not accept connections from these client tools.

- The choice **Version strict applied per study basis** allows the study-level configuration of the same setting to control the setting. This is the less restrictive choice.
- The choice **Only current version supported** specifies that the client and server application versions must match exactly, and it is not possible to override this at the study level. This is the more restrictive choice.
- The choice **Minimum version supported** allows a minimum minor version, within the matching major version, to be specified. Client applications with a version earlier than the one specified here will not be able to connect. All patch numbers, within the matching major and minor version number, are accepted.
- It is only possible to override the Version Strict setting at the study level with a more restrictive setting.
- For further details, also see [Software Version](#) and [Version Strict \(Studies\)](#).
- **Multi-tenant:** Check this box if the server hosts studies for more than one client (tenant). A multi-tenant server places additional restrictions on permission

to use the Image Router and visibility of users that have accounts with other tenants on the server. Enabling this setting is suitable for CROs and organizations that offer **DFdiscover** as a service.

The following **Login Password Security Settings** may be specified:

- **Minimum password length:** Users may change their password to a string no shorter than the value specified here, and no longer than the fixed maximum of 64 characters.
- **Passwords require:** To increase the strength of user passwords you may require that they must contain any combination of the following attributes:
 - at least one character and one digit
 - at least one lowercase and one uppercase character
 - at least one special character (@!#\$)
- **Password expires every:** Passwords may be set to expire every 1-9999 days, after which the user will be prompted to enter a new password on the next login.
- **Disallow reuse of last:** When resetting their password users can be prevented from reusing their previous 1-10 passwords. A value of 0 in this widget allows users to re-enter their current password. This applies to each user's initial password and to any reset passwords, so to prevent users from reusing these passwords after first login set this parameter to at least 1.
- **Lockout users for:** If an incorrect password is entered several times you may suspect that someone is trying to break into your system using another user's login account. To minimize this risk, you can lock the login account for 1-999999 minutes after 1-9999 failed login attempts. A compromise, for forgetful users, might be to lock the account for a short period of time, during which they can look up the password and try again. When a lockout occurs, this message appears at the bottom of the login dialog: *Account locked because of repeated password errors - You may try again in XX minute(s)*.
- **Notification Mail Recipient:** When a login account is locked out a message is written to the system log. In addition, an email can be sent to the addresses specified in this widget.
- **Password reset option:** Selecting this option will allow users to reset their passwords to a temporary single-use password from the login dialog. By default, this option is not selected. The user's email address used for this purpose must be registered under the **Users** tab in the Email field of the contact info section.

A banner may be specified to customize the login dialogs for **DFexplore**, **DFsetup**, **DFadmin** and **DFsend**. The banner appears at the bottom of the login dialog when the user selects the **DFdiscover** server. The banner may consist of one or more lines, include UNICODE characters, and use HTML to add formatting and links to URLs.

Links to media, such as image files are not supported. However, images may be included using base64 encoding (which represents binary data in an ASCII string format). Base64 image encoders are readily available online.

By default, the login banner consists of the organization name, country, state, locality and email contact information.

Studies

A **DFdiscover** study can be thought of as a database of related CRFs. An actual clinical trial might be logically divided into more than one database and thus may be defined as more than one study to **DFdiscover**. For example, while subject CRFs might all go into the main study database, a separate investigator database might be kept for investigator registration forms, an investigator fact sheet, and correspondence with investigators. The key structural difference between these 2 databases is that while a subject ID would serve as the ID key field in the subject database, an investigator number would serve as the ID key field in the investigator database.

In very large studies with many sites it might also be desirable to divide the sites into regional databases, to keep the individual study databases smaller and thereby improve system performance.

In most cases, clinical trial databases are paired as a development database and a production database. Work is done in the development database, tested and then pushed to the linked production database. The trial is conducted with, and data is collected in, the production database.

Study Numbers

A maximum of 999 active studies (or databases) can be defined per **DFdiscover** installation.

NOTE: DF/Net Research, Inc. occasionally makes test and demonstration studies available for use by clients. Since these studies use study numbers in the range 250 through 255, it is recommended that the maximum study number used actually be 249. On an appropriately licensed system, it is further possible to use study numbers 256 through 999 for EDC studies.

Each **DFdiscover** study is assigned a unique number in the range 1 to 999. **DFdiscover** is able to automatically identify and route CRFs that have a **DFdiscover** barcode. For such studies, the maximum study number is 255 (this is the maximum that can be barcoded). If the study will contain paper forms that are submitted for data entry, and should be automatically routed, then the study numbers must be in this range of 1-255. If the study is purely EDC with no scanned CRF images, or if automatic routing is not needed, then the full range of 1-999 can be used.

A study number is the first key necessary to uniquely identify a CRF and route it to the correct database. Because of this study numbers must be chosen carefully, and their uniqueness maintained for the duration of the study. Consideration should also be given to how the numbers are planned. For example, are numbers assigned sequentially starting with study number 1? Or is some other stratification desired, with study ranges allocated for product lines, phase, region, or other natural ways to categorize studies?

Study Directories

A study database is implemented as a collection of files that reside on one or more storage devices. Typically, the collection of files that are a study database are all stored under one unique parent directory. Placing all the files for a particular study under a parent directory that is unique to that study is a good way to organize those files. This eliminates the possibility of mistakenly sharing one file between two or more studies, and makes it easier to encapsulate a study for migration or backup purposes.

When each study has its own parent directory, **DFdiscover** names all sub-directories and study files consistently across studies. Hence the name and location of a file (relative to the parent study directory) will be the same regardless of which study it belongs to.

The default **DFdiscover** file structure for a study includes the following:

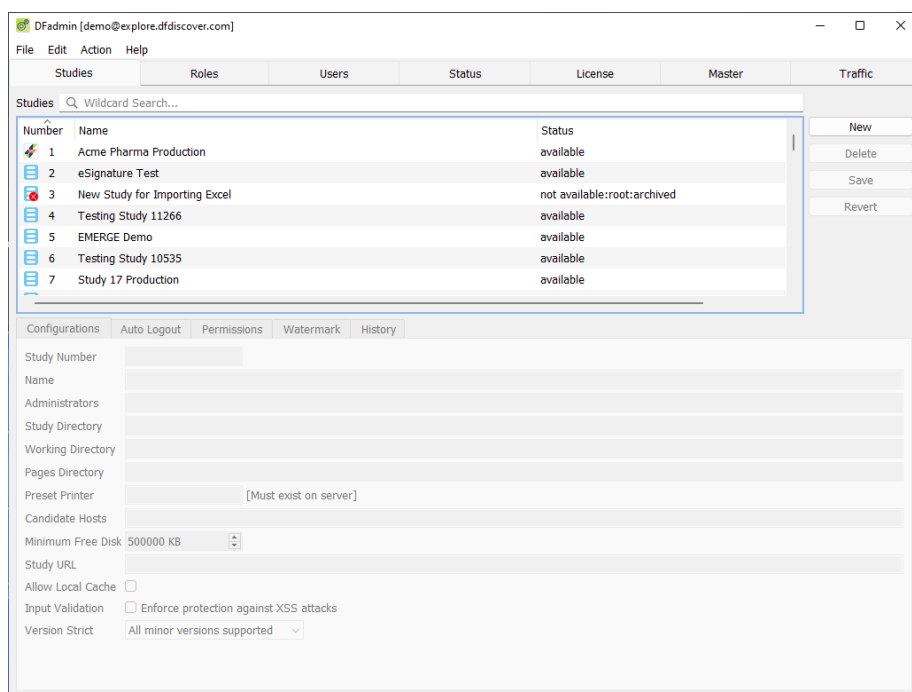
batch	Control files and output logs for study-level batch edit checks.
bkgd	Background image files used by DFexplore , DFsetup , and DFprintdb applications.
data	All DFdiscover database files, including one data file and one index file for each unique plate. It is strongly recommended that this data directory be local to the machine that executes the study database server.
dde	Double data entry sets that are waiting to be processed.
dfsas	Any stored SAS@ jobs that were created from DFexplore .
drf	Any DRFs (Data Retrieval File) created by server-side applications, including reports, and DFexplore .
ecbin	Any study specific scripts or applications that are called from within edit checks using the <i>dfexecute</i> function, or on arrival of a specified plate defined under Plate Arrival Triggers in the DFsetup plates view.
ecsrc	All edit check source files that are specific to this study.
lib	Specific library files, including study server configuration file, sites database, visit map, data dictionary, ICR template.
lut	Any study specific lookup tables.
pages	Base directory for standard resolution (100dpi) CRF pages. Each data record has a corresponding CRF image file below this directory.
pages_hd	Base directory for high resolution (300dpi) CRF pages. Each data record has a corresponding CRF image file below this directory.
reports	Study specific scripts and applications that generate output for review by study staff. This directory also contains all previously created Query Reports.
work	Storage location for study temporary files.

All of these directories are created automatically by **DFdiscover** when a new study is defined. In general, the files and directories should not be accessed directly but instead be accessed through one or more **DFdiscover** applications.

A detailed description of the format of study configuration and data files can be found in [Programmer Guide, DFdiscover Study Files](#).

The Studies Tab

Most study administration tasks are performed using the **Studies** tab in **DFadmin**.



All studies currently registered with **DFdiscover** are displayed, one row per study, in the studies list. Each row shows the unique **DFdiscover** study number, the descriptive study name, and the study status. If the study is not available (either disabled or restricted) the status message will include the administrator's username as well as the message from the administrator. The study list is filtered based on the current user's permissions: the **DFdiscover** administrator can see all known studies in the list (no filtering is applied), while a study administrator can only see the studies they have permission to administer.

Using the Studies view, **DFdiscover** administrators can:

- Register and configure a new **DFdiscover** study.
- Delete a study from the **DFdiscover** study registry. This action does not delete anything from the study directory. It makes the study unavailable to **DFdiscover** users, until/unless it is recovered using one of the 'Add' options.
- Delete all study data. This action is typically used to remove all test data before going live. It removes directories: data, pages, pages_hd, drf, work, and reports/QC, and all dfsas/*.*** files (i.e. DFSas job output files) from disk.
- Delete all study data and setup info. This action removes everything from the study directory except the study configuration file *DFserver.cf*.
- Define a watermark that is applied to all printed output generated by users assigned to any of the roles listed in the watermark definition.
- View and export study change history, including changes to roles and user access.

The first study server 'start' message in the **DFdiscover** server_log corresponds to the date the study was created. A log message identifying the study, user, date and time is also appended to server_log for each of the 3 delete operations.

Use **Wildcard Search** to find studies by their number or name. Wildcard matches are not case sensitive and include the following options:

- * - matches zero or more characters
- ? - matches a single character
- [...] - matches the set of characters in square brackets

The following study settings may be specified in the **Configurations** tab:

Tenant	Tenant is an optional, descriptive label for the client/tenant that "owns" this study. This setting is visible only in installations that enable multi-tenant hosting (see Multi-tenant, Master).
Study Number	Study Number is the unique number for the study.
Name	Name is the name of the study.
Administrators	Administrators are the administrators for the study.
Study Directory	Study Directory is the top-level disk folder where the study is stored on the server.
Working Directory	Working Directory is where the temporary data is saved as it is processed.
Pages Directory	Pages Directory is where images are stored. HD pages are stored in a sibling folder with the same prefix plus a *_hd* suffix.
Preset Printer	Preset Printer is the printer where print operations are directed, unless another printer is specified elsewhere as an override.
Candidate Hosts	Candidate Hosts identifies the servers where slave daemons are running. The value must equal, or be a subset of, the value for Bootstrap Slave Hosts in the Master view.
Minimum Free Disk	Minimum Free Disk is the minimum amount of free disk space for the study. The recommended minimum value for this field is 500MB. The maximum value is 20GB.
Study URL	Study URL is ignored in the current release.
Allow Local Cache	Checking Allow Local Cache enables the DFexplore application to cache study setup information (never study data) on the user's local computer. This will shorten start times for repeated uses of the same study.
Input Validation	When checked, Input Validation enables web-based applications to perform input validation as protection against cross-site scripting (XSS) attacks. While providing enhanced security, input validation changes user inputs and is thus disabled by default. When enabled, if data entered by a user is changed in DFweb, they are warned that their input has changed and asked to review the updated input before saving. It is recommended to enable this option for all studies using DFweb. Studies may want to leave this option disabled only if collecting HTML-based content in fields.
Version Strict	Version Strict specifies how DFdiscover handles connections from client applications that do not exactly match the DFdiscover server software version.

The following study settings are specified in the **Auto Logout** tab:

- **Default** is the auto logout interval automatically assigned to users for this study.
- **Maximum** is the maximum auto logout interval that a user can set.

Adding a New Study

Each new study must be registered before it is used in **DFsetup** to create the study database or in **DFadmin** to define study roles and assign them to users. Studies should be registered early in the study-planning phase so that a study number can be allocated, ensuring that the same number cannot be assigned to another study.

1. Click or select >

The **New Study** dialog is displayed.

2. **Study Number:** Each study database must be assigned a unique number in the range 1 to 999 inclusive.
3. **Study Name:** The study name appears in a variety of places including the study selection dialogs, the top of the screen in the data collection tools and **DFsetup**, and in the header of most **DFdiscover** reports. The study name may be no more than 256 characters.
4. **Study Space:** Before any studies can be registered a **UNIX** administrator must specify permitted study spaces, i.e. the **UNIX** pathname(s) where study directories can be created, and provide a label for each space. Study space directories must be created with *datafax* as the owner, *studies* as the group, and have read-write-execute permission for both *datafax* and *studies*. Each space must be unique (case-insensitive) and cannot itself be a study directory, nor can any sub-component of a study space be a study directory.

If multi-tenant is enabled on the server, it is recommended that each tenant be assigned a unique study space. This further helps with the separation of storage and simplifies backup and restore operations.

Once the UNIX directories are created from a command-line, they can be registered as study spaces by running **DFserveradmin** as root on the **DFdiscover** server, and selecting the study spaces tab.

NOTE: No part of a study space can itself be a **DFdiscover** study directory. For example, `/opt/studies` and `/opt/studies/NIH_Studies` could both be study spaces for the following study directories: `/opt/studies/A`, `/opt/studies/B`, `/opt/studies/NIH_Studies/A`, `/opt/studies/NIH_Studies/B`, etc., but `/opt`, `/opt/studies` and `/opt/studies/NIH_Studies` cannot be **DFdiscover** study directories. Study directories cannot include the follow characters: `\` '$; & * < > |` and SPACE

5. **Study Folder:** Specification of the UNIX path to a new study directory is performed in 2 steps. First select a 'Study Space' from the drop-down list of predefined spaces, and then specify a study folder name, which will be appended to the study space to create the UNIX pathname for the study directory. For example, if the user selects study space name 'NIH_Studies' which has been mapped to UNIX directory `/opt/studies1`, and then enters study folder name 'ECIC' the study directory becomes `/opt/studies1/ECIC`. The study folder name must be unique (case-insensitive) within the study space.
6. **Mode:** Select one of 3 modes:
 - **New Study** Use this mode to create a new **DFdiscover** study.
 - **Restore previous Study at this location** A study which has been deleted using 'Delete-Remove Study from **DFdiscover**' remains on the **DFdiscover** server and can be restored using this option, if you remember the study number, space and folder names.
 - **Restore previous Study with new Study Number** This option is similar to the second option, but allows you to restore the study using a different study number. It does not update any existing data records or barcodes with the new study number, so this is of limited value, but may be useful in some cases. For example, you might have a StudySpace/Test pathname which you clean out using 'Delete-All Study Data and Setup Info' when no longer needed, and restore when you want to use it to setup and test a new study.

Each of these modes has specific requirements and will produce an error message if they are not met. A 'New' study must have a new study number and a space/folder path that does not already exist on the **DFdiscover** server. A study can only be restored using the second mode if the space/folder path exists on the server and the configuration file at that location has the specified study number. And for the third mode the space/folder path must exist, but the study number must be different.

7. **Click **OK**.** This does 3 things on the **DFdiscover** server:
 - The study directory is created (if needed) and the configuration is saved to file `lib/DFserver.cf` in the study path.
 - The study is added to the **DFdiscover** study registry.
 - The study database server starts for the first time, verifies that **DFdiscover** is able to create and write to the study directory, and then creates and initializes the study directories and files. If the file system creation completes successfully, the study list is updated and the focus shifts to the new study entry.

Deleting a Study

Three types of delete can be performed in the **DFadmin** Studies tab, as described below. These options are available, to **DFdiscover** administrators only, by clicking **Delete** or selecting **Action** > **Delete**.

- **Remove Study From DFdiscover...:** This option can be launched from **Delete** as well as from the **Action** menu. It removes the study from the **DFdiscover** study registry, file , and removes study roles and user role assignments from the **DFdiscover** user database, file `DFuserdb.log`, but does not remove anything in the study directory from disk. As a result: the study server can no longer start, the study will no longer appear in study selection dialogs,

the study will not be available to command-line applications, and if any barcoded pages arrive for the study they will be sent to the router.

As long as the study directory is not removed from disk by a **UNIX** administrator the study can be brought back online at any time using [Action](#) > [New](#) and re-entering the study number and directory. If the study roles and user role assignments have been saved they can then be imported, effectively restoring the study to its previous state.

- **Delete All Study Data...**: This option is typically used to clean out test data, images and work files before going live with a new study database. It deletes study directories: *data*, *pages*, *pages_hd*, *drf*, *work*, *reports/QC* and also deletes **DFsas** output files from directory *dfsas*, but keeps any **DFsas** job files.
- **Delete All Study Data and Setup Info...**: This option keeps a study registered in the **DFdiscover** study registry but deletes everything about it, including all data and setup information. The study directory remains on disk but the only thing it contains is the study *lib* directory and file *lib/DFserver.cf*. This option is used to restart a study setup from scratch, or to re-use the study number and directory for a new study.

CAUTION: These options make permanent changes and must be used with great care. They are only available to users who have been granted 'DFdiscover Administrator' status under the 'Permissions' section of the 'User' tab. When one of these options is selected, **DFadmin** presents the user with a confirmation dialog and also requires the user to enter their password before the action is performed. If a mistake is made and data is unintentionally deleted, the only solution is to recover it from a backup.

Modifying an Existing Study

Once a study is added its study number and directory cannot be changed, but other configuration parameters can be modified at any time by either a study or **DFdiscover** administrator.

NOTE: A disabled study cannot be re-configured using **DFadmin** until it is re-enabled, with the exception of the value for **Candidate Hosts**.

The values for the following parameters can be safely changed at any time:

- **Tenant**: This descriptive label is displayed in the Studies tab only and is used solely for internal documentation purposes.
Study Name: A change to the descriptive study name is reflected in any application that starts after the change is saved.
- **Preset Printer**: A change to the default printer takes effect immediately. The printer queue must exist on the server and be listed in */opt/dfdiscover/lib/DFprinters* if one exists, or the entry will not be used.
- **Candidate Hosts**: If the value for **Candidate Hosts** is changed and the study server is already running, the study server will continue to run on the current host. The new hostname(s) will only be considered as possible candidates for the next start of the study server, after the current instance terminates.
- **Minimum Free Disk (KB)**: A change to the value of **Minimum Free Disk (KB)** does take effect immediately.
- **Study URL**: This parameter is ignored for the current release. Although a value can be specified and changes will be saved, the value is not used within **DFdiscover** and will be removed in a future release.
- **Version Strict**: When a client application connects to a **DFdiscover** study server the client version number is compared with the server version number. The result of that comparison, combined with the value of this setting, determines whether or not the connection proceeds. For additional details regarding version numbers, see [Software Version](#).

NOTE: Due to important security improvements, all non-API client connections which do not support TLS v1.3 will be unable to connect to DFdiscover server version 5.8 and later. This includes DFdiscover 5.1 desktop tools. Even when the Version Strict setting allows for 5.1 clients, DFdiscover server 5.8 and later will not accept connections from these client tools.

It is possible for this value to be set at the master level (see [Version Strict \(Master\)](#)), in which case the study setting is only enforced if it is more restrictive than the master setting. A warning will appear if the study setting will be overridden by the master setting.

The choice **All minor versions supported** requires that only the major version number of the client match with the server. Clients with any version or patch number, within the matching major version number, are accepted. This is the least restrictive choice.

The choice **Only current version supported** requires that the major and minor version numbers of the client match with the server. All patch numbers, within the matching major and minor version number, are accepted. This is the most restrictive choice.

The choice **Minimum version supported** allows a minimum minor version, within the matching major version, to be specified. Client applications with a version earlier than the one specified here will not be able to connect. All patch numbers, within the matching major and minor version number, are accepted.

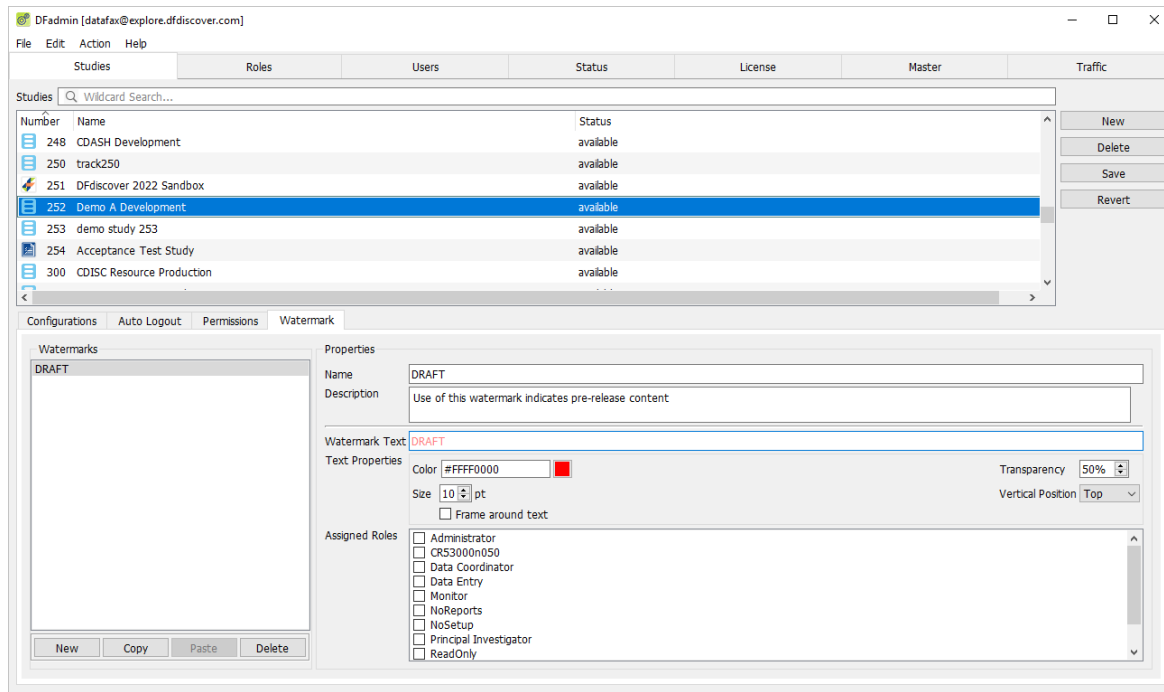
The appropriate choice may depend upon the software changes in minor versions as well as study requirements.

- **Auto Logout - Default and Maximum (Minutes)**: **DFexplore DFweb** and **DFsetup** sessions automatically logout after a specified period of inactivity. In **DFexplore** users can set their auto logout preference to any value between 1 minute and the specified maximum. **DFweb** users can set their auto logout preference to 10, 20, 30, 45 or 60 minutes. The default specification is the initial value for new users. The maximum may not exceed 1440 minutes (24 hrs). Default and maximum auto logout values can also be specified for each study role, where they either inherit the study level values entered here, or override them with values for users with that particular role. Changes to auto logout specifications take effect on a user's next connection to the study server.

Watermarks

Watermarks are used in **DFexplore**, when users assigned to specified roles print PDFs using **DFexplore**. Any pages output by these users are overlaid by the watermark, except for the printing of blank CRF pages.

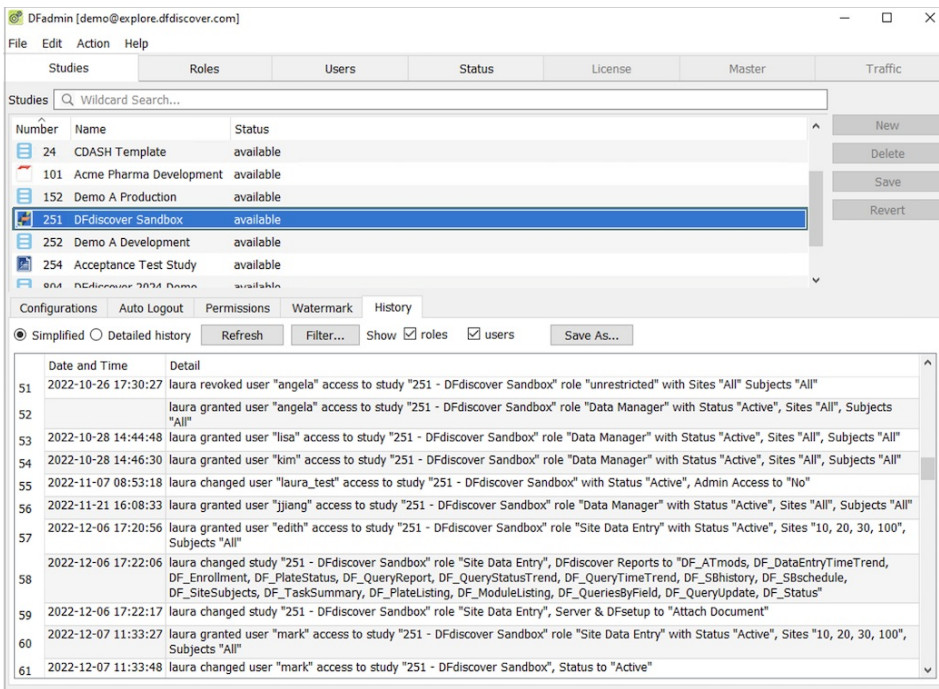
Watermarks are defined in the **Watermark** view.



To define a new watermark:

1. Click **New**. Give the watermark a unique name (maximum 32 characters) or accept the default name provided.
2. **Description** - Provide a description for your watermark (optional - maximum 200 characters).
3. **Watermark Text** - Provide the text of the watermark (maximum 100 characters). UNICODE characters are allowed. Use `\n` to separate multiple lines of watermark text. Watermark text may contain meta-words `%DFUSER` (replace with login user name), `%DFDATE` (replace with current date in format yyyy/mm/dd) and `%DFTIME` (replace with current time in format hh:mm:ss).
4. **Text Properties** - Several properties of the text can be specified:
 - **Color**: Click the color square to display the color chooser dialog, or enter a hexadecimal value directly.
 - **Transparency**: Transparency ranges from 0% to 100%. A watermark with transparency set to 100% is not visible!
 - **Size**: Point size, for the text font, ranges from 6 to 36 points.
 - **Vertical Position**: Select **Top**, **Center** or **Bottom**.
 - **Frame around text**: Check this box to draw a frame around the watermark text.
5. **Assigned Roles** - Check the role names that use this watermark.

Study History



The History tab displays the history of changes to roles and user access for the selected study. User access additions, modifications, and removals performed in the **Users** tab (including study administrator access) and role creation, modification, and deletion performed in the **Roles** tab are shown here. The history displays the date and time of the change, the username of the person making the change, and the details of what was changed. Study history can be exported to an external file in Excel, PDF, HTML, or CSV format.

NOTE: When upgrading to **DFdiscover** Version 5.8.0 or later, existing user and role history must be loaded into the new admin history database for each study, using **DFadmindb** (see [Programmer Guide, DFadmindb](#) for details). Until the **DFadmindb** utility is run, it is not possible to view User, Role, or Study History in **DFadmin**.

The default view is **Simplified**, where changes to a specific role or user are grouped by date and time and described in a single statement. Select **Detailed history** to view the changes outlined in multiple rows and columns according to the specifics of the changes made. The detailed view is designed for easier use in Excel (filtering and sorting).

Study history shows changes starting from when the first study role was created. The current study name is always shown, even if the study name was changed, since study configuration changes are not tracked. Changes are sorted from earliest to latest by default. Click the **Date and Time** column header in the table to reverse the sort order.

Additional actions available in the History tab include the following:

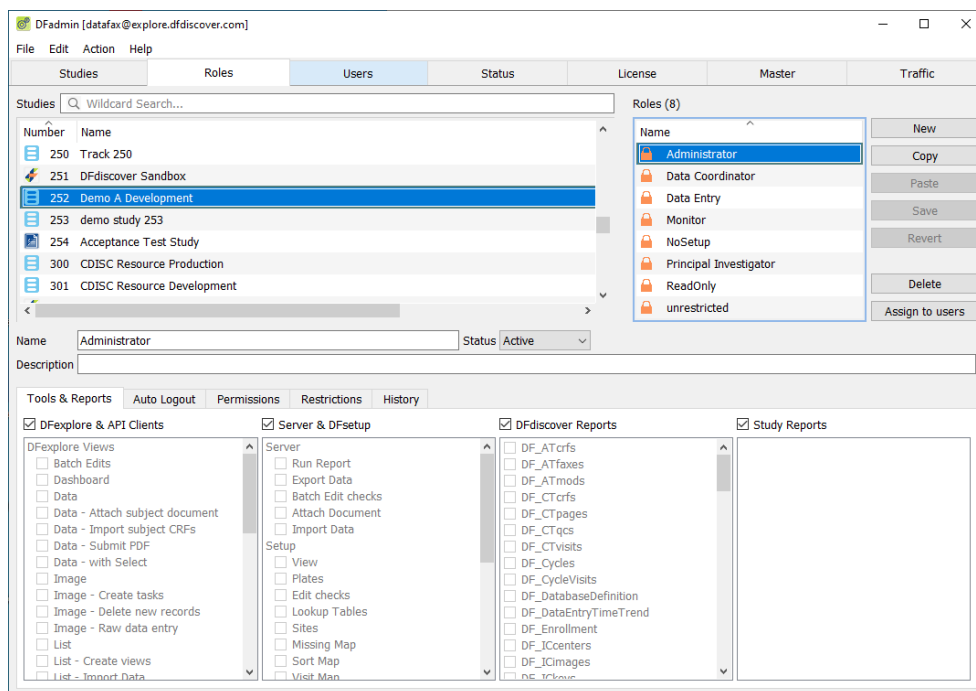
- **Refresh:** See the most recent changes since the table was loaded.
- **Filter...:** View changes made only to specific role(s) by name or to specific user(s) by email, affiliation, or country, as defined in their user profile.
- **Show roles/users:** Hide role or user changes by unchecking the appropriate box. Both boxes are checked by default.
- **Save As...:** Export the current History view to Excel, PDF, HTML, or CSV format.

Roles

User permissions to access studies, perform data entry, run reports, etc. are specified by defining study roles. Only **DFdiscover** and study administrators can create study roles and assign them to users (see [Users](#)). This chapter describes how study roles are created and the full array of specific permissions they encompass.

The Roles Tab

Select the Roles tab.



Roles belong to studies. Selecting a study from the **Studies** list reveals the current roles for that study in the **Roles** list. Use **Wildcard Search** to find studies by their number or name. In this example 10 roles for study number 252 are listed. Role names are case-sensitive and must be unique within a study.

While roles must be defined separately for each study they can be easily copied and pasted from one study to another. Any changes made to role specifications apply only to the study within which the changes are made. Thus, it is possible for 2 roles to have the same name but different permissions in different studies.

The **Status** drop-down is used to switch any role between active and inactive states at any time, without losing the role definition. Users who are granted many roles within a study have permissions associated with their active roles, and if all roles have been made inactive, they are not able to use the study at all.

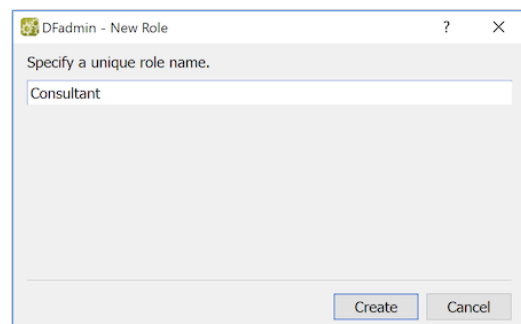
New	Add a new role to the selected study.
Copy	Copy the selected role into memory.
Paste	Paste the most recently copied role to the current study.
Save	Save a new or modified role in the current study.
Revert	Undo all modifications to the current role.
Delete	Delete the selected role from the current study. Deleting a role also removes the role from each user's role assignment list. If a role is deleted and then redefined it will be necessary to reassign it to all users who need it.
Assign to users	Grant users access to a role.

IMPORTANT: When do role changes take effect? Roles can be added, modified and switched between active and inactive at any time, including while users are logged in and working in the study. A user's permission for all studies is evaluated on login to **DFexplore** and thus the user's **DFexplore** instance must be restarted for any changes to take effect.

Adding a New Role

To add a new role to a study, select the study and click **New**. Study status must be **available** before roles can be added, edited or deleted.

Enter a unique role name and click **Create**. Each role name in a study must be unique (case sensitive).



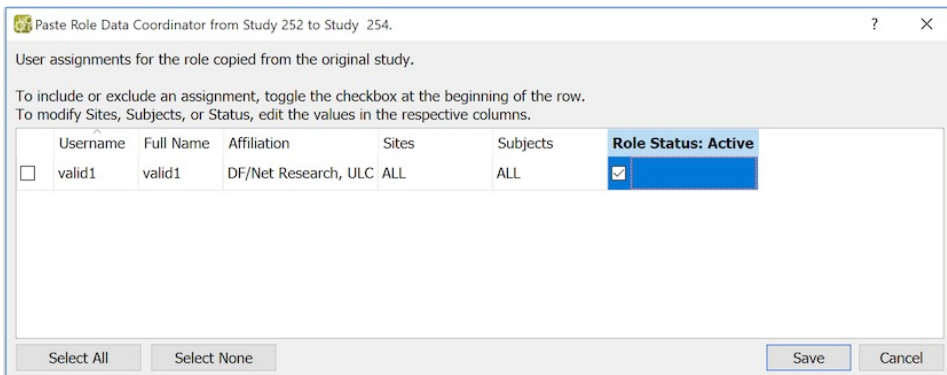
The new role name is added to the study roles list, but the role has no permissions and has not been assigned to any users.

Copying and Pasting a Role

Another way to add a new role is to copy and paste one from the same or another study.

Select a role and click **Copy** to copy it to the clipboard. Paste it any number of times into the same or different studies by selecting the study and clicking **Paste**. Each time a role is pasted, you are prompted to specify a unique role name. If a role with the same name already exists you are required to enter a new, unique name for the role being pasted.

If any users have been assigned the copied role the following dialog will appear when it is pasted.



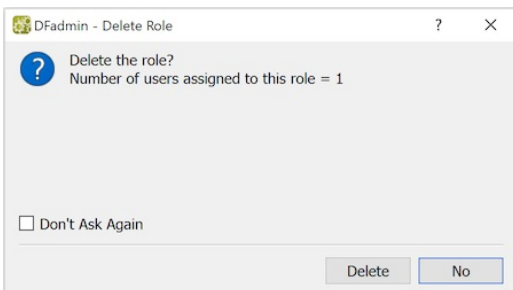
The dialog lists all users who have the copied role, and is used to apply these same role assignments, or modifications of them, when saving the new role. The checkbox in the first column beside each user must be checked to assign the pasted role to the user. By default, this checkbox is off for all users. **Select All** and **Select None** can be used to check and uncheck this box for all users. It is also possible to edit the **Sites**, **Subjects** and **Role Status** specifications for each user before saving the role and user role assignments.

Click **Save** to add the role to the study along with all of the selected user role assignments. Click **Cancel** to cancel both the user role assignments and the addition of the role to the study.

If any of the users could not be assigned the role, due to overlapping role permissions or because the user was locked by another administrator, the role and user role assignments will be saved with the exception of rejected user role assignments (if any).

Deleting a Role

A role can be deleted at any time by selecting **Delete**. When a role is deleted, the confirmation dialog displays the number of users currently assigned the role.



A role can be deleted while users with that role are working in the data collection tool. They will not be affected until the next time they login.

Deleting a role deletes both the role definition and user assignments. If a role is deleted and then re-defined it will be necessary to reassign it to all users who need it.

If you need to temporarily disable a role change the role status to inactive rather than deleting it.

Tools & Reports

This section of the Roles tab is used to specify permissions for **DFexplore** & API Clients, **DFsetup**, some study server operations, standard **DFdiscover** reports and custom study reports. If the box at the top of each of these 4 sections is checked permission is granted to all items in that section. When the same box is unchecked permission is granted by checking each item individually.

NOTE: Using a right mouse click or the keyboard accelerators **Ctrl+A** or **Ctrl+Shift+A**, all items in the current section can be checked or unchecked simultaneously.

The meaning of the individual items is described below.

DFexplore Views

The **DFexplore** View menu is the gateway to the main functional modules that make up this application. Some users may only need access to a few of these

modules. The following permissions grant access to each **DFExplore** view and to specific features within each view.

Batch Edits	Create or retrieve, run and review output from batch edit checks
Dashboard	Enable the dashboard. Also enables for DFweb and DFcollect
Data	Perform data management operations using tasks and subject binders. Also enables for DFweb and DFcollect
Data - Attach subject document	Add a document to associate with a particular subject data record
Data - Import subject CRFs	Allow user to import a PDF file of scanned CRF images
Data - Submit PDF	Allow user to make a secure transmission of scanned CRFs and source documents to the DFdiscover server, as an alternative to emailing or DFsend
Data - with Select	<p>Allow access to the following options under Select in Data View:</p> <ul style="list-style-type: none"> • By Data Field: allow user to build a task set using data field criteria • By Data Retrieval File: allow user to build a task set using an existing data retrieval file • Define Tasks: allow user to define new tasks for themselves and others • Export Tasks to Local File: allow user to export tasks to a local file. • Import Tasks from Local File: allow user to import tasks from a local file. • Change Mode & Level: allow user to change working mode and the workflow level to which records are moved when saved • Batch Validate: allow user to move records in the current task set to a specified level • Batch Sign: allow user to sign records in the current task set to a specified level <p>See additional details about this permission below.</p>
Image	Enter data from CRFs received by email, DFsend , DFExplore Submit PDF, or fax
Image - Create tasks	Allow user to create new data entry tasks within Image View for specified users and study roles
Image - Delete new records	Allow user to delete records in Image View that have not yet been saved to the study database
Image - Raw data entry	Allow user to enter a specified number of new records for specified visits or plates using Image View, but without images
List	Spreadsheet viewer, data exports, DFsas job exports
List - Create views	Create and save record and field selection criteria in named views
List - Import Data	Import data from external files
Queries	Review data queries. Also enables for DFweb and DFcollect
Reasons	Review reasons added to explain data values.
Reports	<p>Run DFdiscover and custom study reports. Only those reports enabled in the DFdiscover Reports and Study Reports section will be available. Also enables for DFweb</p> <p>NOTE: DFdiscover report DF_SBhistory must be enabled in order to view history of changes (audit trail) for plates and fields.</p>
Reports - Create lists	Save a list of reports and report options for subsequent execution
Schedule	Review the current site/subject visit schedule and scheduling requirements
Status	Review current status and workflow level of all records in the study database

NOTE: Data - with Select Permission

In an EDC study this permission should be granted to data management staff only. Clinical sites should *not* have this permission. In addition to restricting access to the Select menu permissions described above, users without this permission also have the following restrictions:

- They can save new data records to level 0 with status pending or to level 1 with status final or incomplete.
- They can revise and save pending level 0 records which remain at level 0 if resaved with status pending, and advance to level 1 if saved with status final or incomplete.
- Once records have reached level 1 or higher these users can no longer save with status pending; they must use final or incomplete.
- The Select menu for such users displays just 2 items: 'All Records' which allows users to select and modify any record for which they have permission, and 'By Task' which allows users to review and modify records which meet pre-specified selection criteria.
- When 'All Records' is elected these users can only save records they have modified, and all such records move to level 1.
- When performing a task, the task definition determines which records are retrieved, whether they can be saved without modification and the workflow level to which saved records are moved, but tasks must be carefully defined as users can only see, modify and save records for which they have the necessary get, modify and write level permissions.

DFexplore Print/Save

Permissions to print and save CRFs, data and report output are under the following permission controls:

Blank CRFs	Print and save blank CRFs in Data view
Completed CRFs	Print and save CRFs containing data values in Data view
Data	Print and save data and metadata records in List, Queries and Reasons views. Print and save database statistics in Status View.
DFsas jobs	Create, modify and export DFsas jobs in List View
Images	Print and save images in Data view, Image View and List View. Also enables for DFweb and DFcollect
Reports	Print and save reports in Reports View. Also enables for DFweb

DFexplore Miscellaneous

Developer	Reload edit checks and lookup tables, and trace edit check execution
May disable edit checks	Permission to turn off edit check execution while reviewing records and performing data entry
Show query status detail	Display internal notes and complete status labels in: Queries View, the Query window within Data view, and in the Queries & Reasons review and approval dialogs.

API Clients

Permission to access specific API client applications are controlled by the items in this section. Supported **DFdiscover** applications include:

DFcollect	Mobile app for online and offline data collection. This option is disabled if the server license does not include the DFCOLLECT feature.
DFengage	Mobile app for collection of electronic patient-reported outcomes. This option is disabled if the server license does not include the DFENGAGE feature.
DFsources	Mobile app for source document submission, including redaction and certification.
DFweb	Browser-based web app for data collection.

NOTE: there is a separate permission for attaching documents in the API client applications, under Server - Attach Document.

Server

Run Report	Permission to run DFreport
Export Data	Permission to run DFexport.rpc and/or DFexport
Batch Edit checks	Permission to run DFbatch
Attach Document	Permission to run DFattach and to attach documents in API client applications.
Import Data	Permission to run DFimport.rpc

Setup

Permission to use **DFsetup** to view study setup specifications or be able to create and modify them is granted as follows:

Status	Use the check box at the beginning of each row to activate or deactivate the database permissions defined in the row. When the box is checked the permissions are active, i.e. included in the role definition and applied to users with that role.
Get Visits	The entry in this cell determines the visits/assessments to which the permissions apply. One or more visits in the range 0 to 65535 may be specified, like this: 0,2,5-10,101 . ALL or * (asterisk) may be entered to represent all visit/assessment numbers.
Get Plates	The entry in this cell determines the plates to which the permissions apply. One or more plates in the range 1 to 501 may be specified, like this: 1-10,21-29,55,68 . ALL or * (asterisk) may be entered to represent all possible plate numbers.
Get Levels	<p>Get Levels determine which data records a user can see in DFexplore. Data records at other levels will be hidden to users with this role. One or more levels in the range 0 to 7 may be specified, like this: 0-3,5,6. ALL or * (asterisk) may be entered to represent all possible levels, and - (dash) represents no levels. If a user has permission to see a data record they automatically have permission to see all of the queries and reasons attached to that record regardless of the levels of the queries and reasons.</p> <p>New blank pages, new pending records, and records with an overdue visit or missing page query have workflow level 0. Without get level 0, users can view new blank pages, cannot view new pending records, and cannot view overdue visit or missing page queries (the record appears as a new blank page instead).</p>
Show Hidden Fields	Fields defined as hidden in DFsetup can be made available in specified roles. The options include all or none . Hidden fields can be made available on some plates but not on others by defining multiple rows, with the desired visit/plate combinations.
Show Internal Queries	<p>Internal queries are typically not made available to the clinical sites but used to track issues within the central team without sharing them with the sites. However, it may be desirable to reveal internal queries after they have been resolved, thus the following options are available: all, resolved and none. Internal query behavior is impacted by modify query permissions (Queries - M) as follows:</p> <ul style="list-style-type: none"> • When none is selected and the role includes full permissions to modify queries, internal queries are visible in DFexplore but hidden in DFweb and DFcollect. • When resolved is selected and the role includes full permissions to modify queries, internal queries are treated the same as external queries in all three applications. • When none is selected and the role includes no or partial (edit check only) permission to modify queries, internal queries are hidden. With partial permissions, internal queries can be created or modified by edit check but the queries remain hidden. • When resolved is selected and the role includes no or partial (edit check only) permissions to modify queries, internal queries are visible and view only in DFexplore and hidden in DFweb and DFcollect. With partial permissions, internal queries can be created or modified by edit check, but they remain view only in DFexplore and hidden in DFweb and DFcollect. <p>NOTE: The internal queries permission will be updated in a future release for consistency between DFexplore, DFweb, and DFcollect.</p> <p>To make internal queries available on some plates but not on others, multiple rows can be defined with the desired visit/plate combinations.</p>
Modify Levels	<p>Modify Levels determine which data records can be modified in DFexplore. Data records at other levels can be viewed but not changed (provided they are included in the user's 'Get' levels). One or more levels in the range 0 to 7 may be specified, like this: 0-3,5,6. ALL or *(asterisk) may be entered to all possible levels, and -(dash) represents no levels.</p> <p>In DFexplore level 0 is used for both new blank pages and new pending records. Thus, a user with Get permission for level 0 records, who does not have Modify permission for level 0, will be able to see blank and new pending records, but will not be able to change them.</p> <p>Permissions specified for queries and reasons only apply to data records the user is allowed to modify. The level of the query or reason itself is irrelevant. If a data record can be changed, query and reason permissions take effect for all queries and reasons on that data record, regardless of the metadata record levels. And if a data record cannot be changed neither can the metadata on that record, even for queries and reasons that are at modify levels. This dependence on the level of the data record also applies to the auto-resolution of missing value and illegal value queries, as described under Query permissions below.</p>

<p>Write Levels</p>	<p>Write Levels determine the workflow levels at which data, queries and reasons may be saved in the study database. One or more levels in the range 0 to 7 may be specified, like this: 0-3,5,6. ALL or *(asterisk) may be entered to represent all possible levels, and -(dash) represents no levels.</p> <p>Write levels are applied separately within each row of the role definition; and can thus be different for different visit/plate combinations.</p> <p>When a data record is saved in <i>Validate</i> or <i>DDE</i> mode the data and all metadata records (both old and new) are written at the user's current working level, whether the records have been modified or not. When working in <i>Modify</i> mode only modified records are saved at the user's working level; unmodified records remain at their current level. As a result, the data and its associated queries and reasons can move to different levels. When working in <i>Edit</i> mode the level remains unchanged for data records, whether they are modified or not, and also remains unchanged for metadata records that are not modified, but new and modified queries and reasons are saved with the level of the data record.</p> <p>Users at the clinical sites should typically have Write levels 0-1. This allows them to enter new records with status pending (level 0) and ensures that all saves with status final or incomplete are moved to level 1. Alternatively, clinical sites can be restricted to Write level 1 only to prevent them from creating new pending records; forcing all records to be saved at level 1 with status Final or Incomplete.</p> <p>A user restricted to write level 0 would be able to create and modify new pending records, but cannot raise them to status incomplete or final. This might be used if an investigator needs to review and approve data entry performed by others before it is elevated to Incomplete or Final status.</p>
<p>Data</p>	<p>Permission to (C)reate, (M)odify, and (D)elete data records, and to mark records (L)ost (now referred to as Missed) are specified separately by checking the box corresponding to each of these functions. If none of these boxes is checked the role will allow DFexplore and API client users to retrieve the data records specified under the Get specifications, but will not allow them to make any changes to these records.</p> <p>The Modify check box has a middle state, denoted by a dash or shading (depending on platform). This signifies that the user is allowed to change a record's status and level but not modify any of the data fields. This feature is useful for site monitors who need to perform source verification, and perhaps add queries, but not change data values. Typically, such users will use predefined tasks which will retrieve the records that need to be reviewed, and move them to a designated level when they click a save button to indicate that the review was completed. This mode allows data fields to be modified by edit checks that might be used within the task to update hidden fields on each plate that is reviewed with the user, date and time of the review, or other information obtained from the reviewer using edit check functions <i>dfask</i> or <i>dfcapture</i>.</p> <p>Checking the (P)assword required box indicates that the user must provide their eSignature when saving the record. This only applies for plates with an eSignature module where the record is eligible for signing. The user is prompted for both username and password on the first signing, then just password on subsequent signings in a given login session. Once the credentials are accepted, the eSignature fields are automatically populated. Refer to Study Setup User Guide, eSignatures and 21 CFR Part 11 Compliance for further details. This functionality is supported in DFexplore, DFweb, and DFcollect.</p> <p>For the following DFexplore operations, users are always required to enter their password, regardless of role specifications:</p> <ul style="list-style-type: none"> • Delete a page • Batch validate a set of records in Data View • Batch sign a set of records in Data View • Import Subject Documents in Data View • Import Subject Data in List View • Apply data changes when running batch in Batch Edits View <p>If the credentials are not accepted, the user is prompted to re-enter their credentials. If they fail to do so in the configured number of attempts, they are automatically logged out and a notification email is sent to the DFdiscover administrator.</p> <p>It is also possible to require entry of the user's password in an edit check using function <i>dfpassword</i>, for example before: unmasking certain data fields, revealing certain data entry screens, or running a randomization script.</p>

Queries

Query permissions are granted by checking the box associated with (C)reate, (M)odify, (D)elete, (A)pprove and (R)eply. Each of these permissions has three levels: not allowed (unchecked), allowed (checked), and allowed only when performed by an edit check (shaded or dashed).

DFexplore includes a number of rules for automatically changing query status. These rules and the permissions required to enable them are described below.

1. Query status changes from unresolved to pending if a new reply is entered. This rule applies to all queries. Of course the user must have permission to reply to queries on the current record, and the data record must be at a workflow level for which the user has modify permission.
2. Query status changes from unresolved to pending if a new reason is entered, subject to the following conditions:
 - the query usage type is external (not internal)
 - the data record level is included in the user's modify levels
 - the user has permission to modify queries; either full permission (checked) or edit check permission (shaded or dashed).
3. In **DFexplore** queries, with category missing value or illegal value, can be auto-resolved when the user corrects the field, and auto-unresolved if the user changes the field back to a missing or illegal value. This only occurs when all of the following conditions are met:
 - the user has just changed the value of the data field
 - the field has an external query (not internal)
 - the category code is missing or illegal
 - the query reply type is refax/correction
 - the query status is unresolved or resolved (not pending)
 - the data record level is included in the user's modify levels
 - the user has permission to modify queries; either full permission (checked) or edit check permission (shaded or dashed).

When these conditions are met the following rules are applied in order until a rule is met. For Illegal value queries:

- if the category code was illegal at plate entry, and the field has been returned to the value it had at plate entry, then the query status is returned to the status it had at plate entry
- else, if the new value is illegal, then query status is set to new unresolved (code 1)
- else, if the new value is a missing value code, then query status is set to resolved NA (code 3)
- else, the new value must be legal, and thus query status is set to corrected (code 5). Note: if a field does not have legal value specifications all values are considered legal, and if a field is optional, a blank value is considered legal. legal.

For missing value queries:

- if the category code was missing at plate entry, and if the field has been returned to the value it had at plate entry, then the query status is returned to the status it had at plate entry
- else, if the new value is blank or illegal, then query status is set to new unresolved (code 1)
- else, if the new value is a missing value code, then query status is set to resolved NA (code 3)
- else, the new value must be non-blank and legal, and thus query status is set to corrected (code 5). Note: if a field does not have legal value specifications all values are considered legal.

Typically roles assigned to **DFexplore** users at the clinical sites should allow them to: (R)eply to queries (checked), (C)reate, (M)odify and (D)elete queries only when this is performed by an edit check (shaded or dashed), and not allow them to (A)pprove replies to queries (unchecked); while roles assigned to central data management staff should allow them to: (C)reate, (M)odify, (D)elete and (A)pprove queries (checked), but not (R)eply to queries (unchecked). Role assigned to Monitors may have the same query permissions as central data management staff. To approve queries, both (A)pprove and (M)odify permissions must be enabled.

In **DFweb and DFcollect**, in addition to enabling the (R)eply query permission, the user must also have (M)odify permissions to reply to the query. If checked, the user can reply to and also modify the query. If shaded or dashed, the user can reply to but not modify the query except via edit check. If not enabled, the user will encounter an error and will not be able to save the record.


Reasons	<p>Reason permissions are granted by checking the box associated with (C)reate, (M)odify, (D)elete, and (A)pprove. Each of these permissions has three levels: not allowed (unchecked), allowed (checked), and allowed only when performed by an edit check (shaded or dashed).</p> <p>As for queries, reason permissions only take effect on data records the user can modify, and the level of the reason record itself is ignored.</p> <p>Typically roles assigned to data collection tool users at the clinical site should allow them to: (C)reate and (M)odify reasons (checked), but not allow them to (A)pprove reasons (unchecked) unless there are also edit checks that create reasons for them (shaded or dashed); while roles assigned to central data management staff should allow them to: (A)pprove reasons (checked), but perhaps not (C)reate or (M)odify them, except when this is done by an edit check (shaded or dashed).</p> <p>If a user has permission to approve reasons, any reasons they create or modify will be automatically approved, whether created manually or by function <i>dfaddreason</i></p> <p>If users do not have permission to approve reasons any reasons they create or modify will be saved with the pending status, whether created manually or by function <i>dfaddreason</i>.</p> <p>Users with shaded or dashed permission cannot approve reasons themselves, but approved reasons can be created by edit checks. For these users reasons created/modified by <i>dfaddreason</i> will be automatically approved if the edit check adds the reason without user intervention, but if the edit heck displays the reason dialog so the user can approve or modify it then the reason will be saved with status pending.</p> <p>Automatic "Set by Edit check" reasons, which are created when an edit check changes a data field, are automatically approved for all user regardless of reason approval permission, but these reasons can be overridden within the edit check using the <i>dfaddreason</i> function to create a custom reason with a specified status.</p>
----------------	--

NOTE: Effect of Get restrictions on edit checks
Restrictions on records the user can Get (i.e. see), specified by Visit, Plate and Level, have implications for edit checks in **DFexplore**. A user must be able to at least view a data record before edit checks can see it. Any reference to a field the user does not have permission to view returns missing in an edit check. This restriction does not apply to **DFbatch**.


NOTE: Invalid database permissions
A row that is highlighted in red indicates:
- an invalid entry in at least one of the columns, or
- overlapping visit/plate combinations; to avoid permission conflicts each visit/plate combination can appear in only one row in the permissions table.

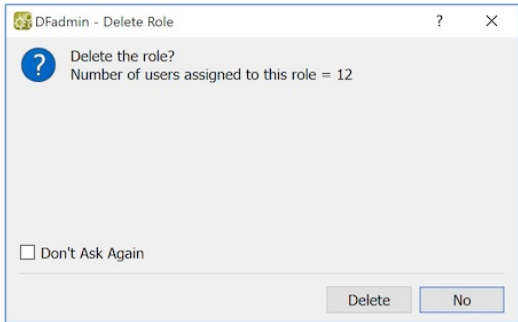
Adding Database Permission

To add a new database permission:

1. Select the first empty row in the permissions table.
2. Enter the visit/plate combination to which the permissions apply.
3. Enter the remaining permission specifications. Press **[Enter]** or **[Tab]** to check the validity of the specifications.
4. If none of the values appear in red, the row will be accepted as valid, the delete icon () will appear in the last column, and a new empty row will be added to the permissions table.
5. When all of the rows needed to define the database permissions for the role have been completed, select **[File]** > **[Save]** to save the permissions. Any changes made to the permissions are not saved until they are saved explicitly in this way.

Deleting Database Permission

To delete a specific database permission specification (table row), click the corresponding delete icon (). A confirmation dialog is displayed, warning you of the number of users who are affected by this modification to the role definition. Click **[Delete]** to confirm and proceed with the deletion; otherwise, click **[No]** to cancel the deletion.



Modifying Existing Study Permissions

Permissions can be changed by clicking and overwriting the value or changing the check box state (checked, shaded/dashed or unchecked), followed by pressing **Enter** or **Tab** to verify that the value is valid.

Rows highlighted in red, indicate errors in role definition. Role errors can occur because an individual value specification is invalid, or because a visit/plate combination overlaps with the combination specified in another row. Invalid specifications cannot be saved. Permissions containing both errors and some new or modified entries that are valid can be saved, but rows containing errors are dropped if they were newly defined, and are returned to their previous valid values if the error was introduced by a modification.

Restrictions

For each role, permissions for specific query categories may be restricted by visit and plate, so that query permissions to (C)reate, (M)odify, (D)elete, (A)pprove, and (R)eply to queries as defined in the Permissions tab may be excluded for specific query categories. These restrictions apply to manual queries only, not to query functions performed by edit check. For details on query categories, refer to [Study Setup User Guide, Query Category Map](#).

The following example shows query category restrictions for an example Monitor role. Their Permissions tab (not shown) allows them to (C)reate, (M)odify, (D)elete, (A)pprove queries on all records. The Restrictions tab (shown) removes their (C)reate and (D)elete permissions for all query category codes except 30 ("Monitor" custom category). They are able to (M)odify and (A)pprove queries of all categories. (R)eply is left blank since they do not have reply permissions.

	Visits	Plates	No C	No M	No D	No A	No R
<input checked="" type="checkbox"/>	ALL	ALL	1-29,31-35		1-29,31-35		
<input checked="" type="checkbox"/>							

Role History

Date and Time	Detail
2018-03-05 17:10:00	1 datafax added role "Monitor" to study "252 - Demo A Development" with Status "Active", Description "Monitor", DFExplore & API Clients "All", DFdiscover Reports "All", Study Reports "All", Server & DFsetup "All", Initial Auto Logout "Inherit from Study", Maximum Auto Logout "Inherit from ...
2018-03-05 17:10:01	2 datafax granted study "252 - Demo A Development" role "Monitor" permission with Status "Active", Visits "All", Plates "All", Data "Create, Modify, Delete, Missed", Queries "Create, Modify, Delete, Approve, Reply", Reasons "Create, Modify, Delete, Approve", Levels "All", Modify Levels "All", Write Levels "All", Hidden Fields "Show all", Internal Queries "Show all"
2024-02-21 12:31:34	3 demo added query restrictions to study "252 - Demo A Development" role "Monitor" with Visits "All" Plates "All", Status "Active", Add "1-29,31-40", Edit "1-29,31-40", Delete "1-29,31-40", Approve "1-29,31-40"
2024-02-21 12:31:53	4 demo changed study "252 - Demo A Development" role "Monitor", cleared Study Reports "All"
2024-02-21 12:32:06	5 demo changed study "252 - Demo A Development" role "Monitor", Server & DFsetup to "Attach Document"

The History tab displays the history of changes to the selected role, including creation, modification, and deletion. The history displays the date and time of the change, the username of the person making the change, and the details of what was changed. History can be exported to an external file in Excel, PDF, HTML, or CSV format.

NOTE: When upgrading to **DFdiscover** Version 5.8.0 or later, existing user and role history must be loaded into the new admin history database for each study, using **DFadmindb** (see [Programmer Guide, DFadmindb](#) for details). Until the **DFadmindb** utility is run, it is not possible to view User, Role, or Study History in **DFadmin**.

The default view is **Simplified**, where changes are grouped by date and time and described in a single statement. Select **Detailed history** to view the changes outlined in multiple rows and columns according to the specifics of the changes made. The detailed view is designed for easier use in Excel (filtering and sorting).

Changes are sorted from earliest to latest. Click the **Date and Time** column header in the table to reverse the sort order.

Additional actions available in the History tab include the following:

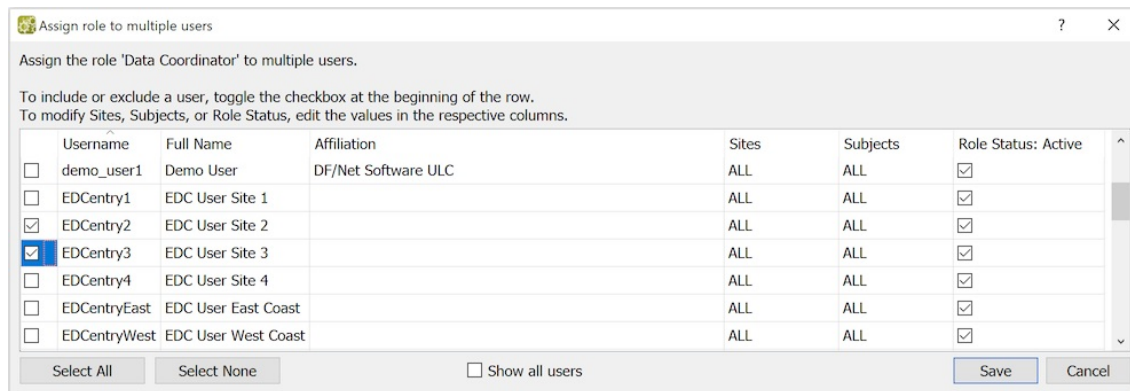
- **Refresh**: See the most recent changes since the table was loaded.
- **Save As...**: Export the current History view to Excel, PDF, HTML, or CSV format.

Assign Roles to Users

A role can be assigned to any or all users by selecting the role and clicking **Assign to users**. A confirmation dialog is displayed listing all users currently using this study, their full name, affiliation and user settings. You may modify sites, subjects or role status from this dialog.

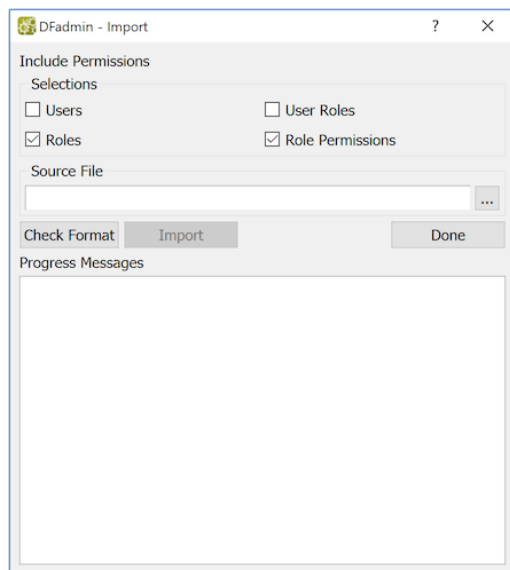
If you want to see a list of all users, mark **Show all users**.

Select the users to receive this role assignment (or select all using **Select All**) and click **Save**.



Import

Study roles, role permissions, role restrictions, role assignments to users, and user contact information can be imported from a text file. Import will overwrite any matching specifications that already exist. The import dialog includes options for selecting the type of records to be imported (i.e. Users, User Roles, Roles and Role Permissions and Restrictions) from the import file as illustrated below. When **Role Permissions** is selected, both Role Permissions and Role Restrictions are imported. The import file must be formatted as described in [DFuserdb.log](#), excluding the Record Time Stamp and Modifier fields.

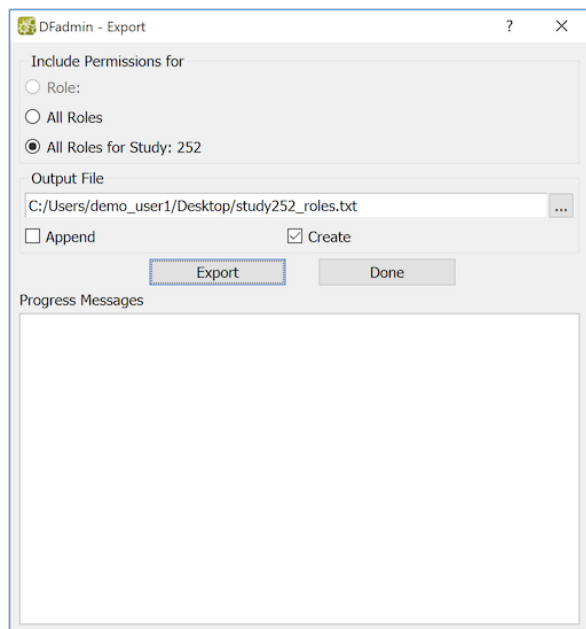


Importing Users, Roles and Permissions

1. Select **File** > **Import**
2. Specify the types of data to import. This is specified by checking or unchecking the items in the **Selections** section of the dialog.
3. Specify the **Source File** name or pick a file by clicking **...**.
4. Click **Check Format**. Any invalid entries are displayed. **Import** is enabled only if there are no invalid records.
5. Click **Import**. Progress and errors messages (such as overlapping permissions) are displayed while importing.
6. Click **Done** to dismiss the dialog.

Export

Study roles, role permissions, role restrictions, role assignments to users, and user contact information can be exported to a text file. Export can overwrite or append its output to the specified file. The output file format is described in [DFuserdb.log](#).



Exporting Users, Roles and Permissions

1. Select the study, user or role to be exported (if export is to be limited to a single user, study or role).
2. Select **File** > **Export**.
3. Select the users, studies or roles to be exported (all or currently selected study, user or role).
4. Specify an **Output File** name or pick the file using the file selection button.
5. Select **Append** to write the exported data to the end of file, or **Create** to overwrite the file with the exported data.
6. Click **Export**. Progress and errors messages are displayed while exporting.
7. Click **Done** to close the Export dialog.

Users

Users must be granted permission to access **DFdiscover** studies. A **DFdiscover** user account is created using the **Users** tab in **DFAdmin**, and involves entering user contact information, preferences and study permissions. This account is required for both **DFdiscover** and the data collection tools. In addition, a UNIX user account is required for **DFdiscover**, but not for **DFexplore**.

DFdiscover user accounts can be modified by any **DFdiscover** administrator or study administrator, but there are limitations.

- Only a **DFdiscover** administrator can grant and revoke **DFdiscover** system and study administrative privileges for other users.
- A study administrator can create and update study roles and **DFdiscover** user accounts, and can grant and revoke study roles to any user in the **DFdiscover** user database. Study administrators are restricted to performing these functions in their own studies.

To administer user permissions, two applications are available: the **Users** tab in **DFAdmin** described here and the command-line application [Programmer Guide](#), [DFuserdb](#).

The Users Tab

Click the **Users** tab to display the **Users** view.

DFAdmin [demo@explore.dfdiscover.com]

File Edit Action Help

Studies Roles Users Status License Master Traffic

Users (205) All Administrators

Username	Password Expiry	2FA Status	SSO Only	Full Name	Affiliation	Master
Demo2025	2053-05-15			Demo2025	DFnet	Admin
demo	2053-08-31			Demo Account	DFnet	DFdiscover
demo_DM	2053-04-27			Demo DM Account	DFnet Biometrics	DFdiscover
demo_user1	2025-05-04			Demo User	DF/Net Software ULC	101,252

Username: demo Status: Active New Password: []

User Profile Preferences Permissions ePRO Admin Permissions History

Full Name: Demo Account

Affiliation: DFnet

Address: 4724 52nd Avenue South

City: Seattle

State: Washington Postal Code: 98118

Country: USA

Telephone: 1-206-322-5931

Fax: []

Email: dfdiscover.demo@gmail.com

Two-factor authentication required SSO Only

Send code by: email Azure SMS

Actions: New / Find, Duplicate, Delete, Save, Revert, All Users, Filter by Roles, Wildcard Search, Add Bulk Users

The Users list displays each user's: username, password expiry date, two-factor authentication (2FA) status, Single Sign-On (SSO) permission, full name, affiliation and administrative privileges. **DFdiscover** administrators are shown as **DFdiscover**, study administrators are shown as **All Studies** or the list of study numbers they are allowed to administer. A blank password expiration date indicates that the user has not yet been assigned a password.

The following actions are available.

All	Show all users currently defined in the DFdiscover users database.
Administrators	Show DFdiscover system and study administrators only. The Admin column shows ' DFdiscover ' if the user is a DFdiscover administrator, 'All Studies' if the user has permission to administer all studies, and a list of DFdiscover studies if the user has permission to administer selected studies.
New / Find	Add a new DFdiscover user account or find an existing user account by email address.
Duplicate	Add a new DFdiscover user with the same account details and permissions as an existing user.
Delete	Inactivate a DFdiscover user account and remove it from the list of displayed users. Once defined a user account is never actually deleted; thus users listed in study audit trails can always be identified.
Save	Save any changes made to current user account or user permissions.
All Users	Show users without applying any study role filtering. If 'All' choice is selected (from the All or Administrators choices), all DFdiscover users will appear in the list. Otherwise, only DFdiscover system and study administrators will be displayed.
Filter by Roles	Select this option to filter the users list by study roles.
Wildcard Search	Select this option to display users that contain a specified text string in either the Login, Full Name or Affiliation fields. Wildcard matches are not case sensitive and include the following options: <ul style="list-style-type: none"> • * - matches zero or more characters • ? - matches a single character • [...] - matches the set of characters in square brackets Thus for example: <ul style="list-style-type: none"> • *er - matches usernames, full names and affiliations ending in 'er'. • ????? - matches usernames, full names and affiliations that have exactly 5 characters • [ab]*er - matches usernames, full names and affiliations that begin with 'a' or 'b' and end in 'er'.
Status	[Status] must be set to Active to enable a user login, and may be set to Inactive at any time to disable it. The head icon appears as an empty outline if login is disabled. Inactivating a user login retains all user information, compared to 'Delete' which removes it.
New Password	Select this button to generate a temporary password for the user to use on their initial (or next) login, at which time the user will be prompted to change the temporary password. The system-generated password conforms to the requirements specified under the Master tab. It can be modified by editing the password field, but the value must conform to the password rules. <div style="border: 1px solid #ccc; padding: 5px; margin-bottom: 5px;"> <input type="button" value="New Password"/> <input type="password" value="badpass"/> An invalid password. </div> <div style="border: 1px solid #ccc; padding: 5px;"> <input type="button" value="New Password"/> <input type="password" value="5ty78-R4#"/> A valid password. </div> When a new password is saved, the list of users is updated and the new password expiry date is displayed.
User Profile	Enter user information and 2FA and SSO preferences. See User Profile for details.
Preferences	Record the user's preferred language and contact method (fax or email). This information is not currently used by DFdiscover .
Permissions	Define the roles a user is granted in different studies.
History	View and export the history of changes to the current user.

User Profile

ePRO User Accounts	<p>A user account is considered an ePRO user if it has one or more ePRO data role(s) and no other roles assigned.</p> <p>User profile data for ePRO users is not displayed in DFadmin to protect personally identifying information (PII) of study participants. Instead, a message is displayed indicating that the user profile is hidden. User profile data is accessible to authorized users from the DFweb ePRO Management page.</p> <p>Only the 2FA setting may be enabled on the user profile in DFadmin for ePRO users. The SSO only setting is not supported in DFengage at this time.</p>
Full Name	The user's full name is used to populate the signature name field when executing an electronic signature.
Email	<p>For regular users, the Email address is required for proper DFdiscover operation. Without a valid email address, users will not be able to perform two-factor authentication via email or reset their own password using the Password Reset option in DFsetup and DFsend login dialogs.</p> <p>NOTE: An administrator must enable the system-wide capability for users to reset their own passwords.</p> <p>If the user is an administrator, their email address is also used as the sender identification when performing certain actions, for example when they restrict a study.</p> <p>It is recommended that an email address be provided for each registered user.</p>
Telephone	<p>A valid phone number is required if two-factor authentication is enabled and Send code by Azure SMS is selected on servers where this option is enabled.</p> <p>The phone number must include both the country code and the area code, e.g. 1-555-1234567.</p> <p>The phone number must be within the region supported by the server's Azure SMS configuration. See Azure SMS Messaging for details.</p>
Two-Factor Authentication (2FA)	<p>Two-factor authentication can be enabled as an added level of account security. When enabled, the user will be required to enter a unique 6-digit security code once every 30 days for each device that they use to access DFdiscover. Certain conditions, such as password resets, can reduce this interval.</p> <p>Security codes are sent via email by default. If enabled for the server, security codes may be sent by Azure SMS text message instead of email, based on the Send code by setting in the user profile.</p> <p>If sending by email, a valid email address is required.</p> <p>If sending by Azure SMS, a valid phone number is required (including both country code and area code), and the number must be within a region supported by the server's Azure SMS configuration.</p> <p>Each server supports SMS only to specific regions (for example, North America by default unless otherwise configured). If SMS is enabled for a user outside the supported region, they will not receive their security code and will be unable to log in. See Azure SMS Messaging for details.</p> <p>Changes to the 2FA settings are tracked under the History tab.</p>
SSO Only	<p>Single Sign-On can be assigned by an Administrator, mandating this as the only method of login for the user. When applied the user will be unable to login using their DFdiscover credentials (if previously given).</p> <p>This feature is intended for users who will always log in via SSO (no username or password given). When SSO is enabled, an error message will be displayed when login is attempted with username and password.</p> <p>NOTE: The Forgot Password option will be disabled if SSO only is enabled.</p> <p>Changes to the SSO setting are tracked under the History tab.</p>
Data Encryption	<p>User profile data is encrypted at rest to ensure that anyone with backend access to the user database may not access personally identifying information (PII) without an authorized DFdiscover user account.</p> <p>User profile data is encrypted when saved to the server and decrypted when retrieved by an authorized user or system via DFadmin, DFexplore, DFweb, DFws API, DFuserPerms, or DFuserdb -export.</p>

Adding a New User


Users are not able to access any **DFdiscover** applications until they have been registered as a **DFdiscover** user.



NOTE: A **DFdiscover** username is a unique and permanent user identifier. Once created, a new user account is retained in the **DFdiscover** users database even if the user is deleted. While it is possible to reactivate a deleted user account and modify it for use by a different user, this is not recommended if the user account has been used to process data for any study, as it could result in confusion in the interpretation of audit trail reports. When adding a new user, chose the username carefully and make sure that it has been entered correctly before saving a new user account.


NOTE: The appearance of the **New User** dialog, and the procedure for adding users, is affected by the setting of **Multi-tenant** in the **Master** panel (see [Multi-tenant, Master](#)). In multi-tenant installations, a study admin is permitted to create new user accounts or assign new study permissions to existing user accounts. To locate existing user accounts, the study admin must know the email address of the existing user account.

1. Click **New / Find**. With **Create new user** selected, enter the new user's unique username. Alternatively, with **Find existing user** selected, enter the email address of the user you wish to add. It is not possible to create a new user account starting from an email address. Users are uniquely identified by their username, which must consist of 16 characters maximum from the list: A-Z, a-z, 0-9, underscore. The first character of the username must be from the list:




A-Z, a-z.

Unlike usernames, email addresses are not required to be unique. When entering an email address, click **Find** to locate existing accounts using that email address.



If the email address is defined, select the desired user account from the list.



2. Click **OK** to add the user and exit the dialog.

If the username is associated with a deleted user account, a confirmation dialog opens asking if you wish to re-activate the deleted user account. Click **Yes** to retrieve the deleted user account information including user contact info, preferences and permissions, all of which can then be modified as needed.



Duplicating an Existing User

When creating a new user account that will have study roles and permissions identical or similar to an existing user, it may be easier to duplicate an existing user.

1. Select the user whose permissions will be duplicated. Highlight the user in the **Users** list.
2. Click **Duplicate**. The **Duplicate** dialog appears. The dialog appearance is almost identical to the **Add** dialog.
3. Enter the new user's username.

Users are uniquely identified by their username, which must consist of 16 characters maximum from the list: A-Z, a-z, 0-9, underscore. The first character of the username must be from the list: A-Z, a-z.

4. Click **Create** to create the duplicate user. The new user account is created, with data copied from the duplicated user, and the new username is displayed in the **Users** list. The new user account can then be modified as needed to correct contact information, preferences and permissions.

Deleting an Existing User

Deleting a **DFdiscover** user account changes the account status to deleted, but does not actually remove it from the **DFdiscover** users database. A user is locked when it is selected in **DFdiscover** and thus cannot be deleted if it is currently locked by another **DFdiscover** or study administrator.

1. Select the user to be deleted. Highlight the user in the **Users** list.
2. Click **Delete**. A confirmation dialog appears. Click **OK** to confirm the deletion, thereby deleting the user from the **Users** list and revoking all of their **DFdiscover** permissions, or click **Cancel** to cancel the delete action.

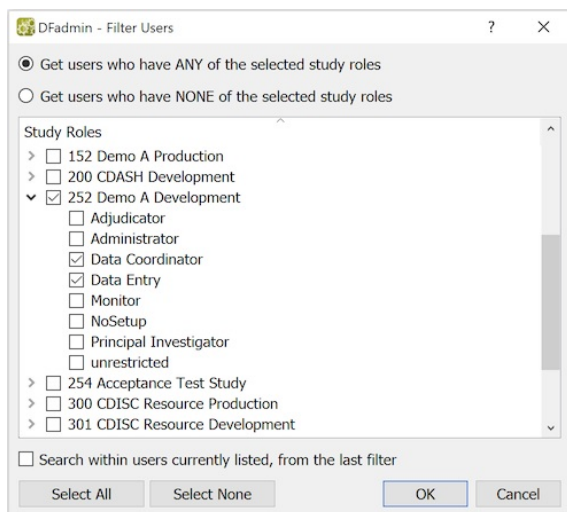
NOTE: Deleting an active user

It is safe to delete a user who is currently logged in and working in a **DFdiscover** study. Users issuing commands from a terminal window will be affected on the very next command, but deleting a user will not have an immediate effect on users of **DFdiscover** applications, including **DFexplore**. User permissions are checked when an application is started. Thereafter, the permissions are not checked again until the application is re-started. An active user will be able to continue working until they exit the application.

Filtering Users

Filtering the user list is an easy way to find users who have specified roles for specified studies. It can also be used to find users who do not have specified roles.

1. Click **Filter by Roles**. The dialog lists all **DFdiscover** studies and the roles for each study. Individual study roles can be selected, or all roles for a study can be selected, using the check boxes beside study and role names. After selecting studies and study roles, use the toggle options at the top of the dialog to indicate whether you want to list users who have any of the selected roles, or conversely list users who have none of the selected roles. After performing an initial search, you can refine it by marking **Search within users currently listed, from the last filter**. If this box is not checked filtering is applied to all registered **DFdiscover** users.



2. Click **OK** to apply the filter. Users who meet the filter specifications are listed in the Users list in the main window. After filtering users, click **All Users** to restore the complete list of registered **DFdiscover** users.

Adding New Users in Bulk

For users submitting data via electronic patient reported outcomes (ePRO) applications such as **DFengage**, their permissions must be set up with a specific study role and access to one subject only. It may be expedient to create accounts for these users all at once at the start of the study. This may be accomplished using the Add Bulk Users function, which creates up to 100 user accounts at a time.

1. With the Users tab selected, click **Add Bulk Users**.
2. The Add Bulk Users dialog appears. Select the study and role and specify the username prefix and the range of subject IDs for the users to be created. Up

to 100 subject IDs can be specified. The subject IDs specified must be defined in the Sites configuration in **DFsetup**. Provide the initial, temporary password to be used for all accounts.

DFadmin - Add Bulk Users

Specify the study, role, username prefix and range of subject IDs to create multiple user accounts each with access to a single subject ID. Each subject user account will be assigned a unique username based on their subject ID and the same temporary password for all accounts.

Study: 117 DFengage eDiary

Role: ePROdata

Username Prefix: s117p

Subjects Range: 10-15,20-25,30-35 Preview usernames

Initial Password:

s117p10-s117p15, s117p20-s117p25, s117p30-s117p35

Add 18 user accounts

Results:

Close

3. Once all information is provided, click **Add X user accounts** to continue, where X is the number of user accounts to be added.
4. Review the details of the new users to be added, including the study number, role name, and usernames. Click **Yes** to proceed with adding the accounts. Click **No** to cancel and return to the previous step.

DFadmin - Users

Add 18 user accounts with the following usernames and grant study '117', role 'ePROdata' permission for the respective subject records.

s117p10-s117p15, s117p20-s117p25, s117p30-s117p35

Click 'Yes' to add the accounts

Yes No

5. The Results panel in the Add Bulk Users dialog will update with the username and status of all the users defined above.

DFadmin - Add Bulk Users

Specify the study, role, username prefix and range of subject IDs to create multiple user accounts each with access to a single subject ID. Each subject user account will be assigned a unique username based on their subject ID and the same temporary password for all accounts.

Study: 117 DFengage eDiary

Role: ePROdata

Username Prefix: s117p

Subjects Range: 10-15,20-25,30-35 Preview usernames

Initial Password:

s117p10-s117p15, s117p20-s117p25, s117p30-s117p35

Add 18 user accounts

Results:

	Username	Log message
1	s117p10	Success
2	s117p11	Success
3	s117p12	Success
4	s117p13	Success
5	s117p14	Success
6	s117p15	Success
7	s117p20	Success
8	s117p21	Success

Close

Permissions

User permissions come from the roles they are assigned within each study. The creation of study roles is described in [Roles](#). This section describes how study roles are assigned to users.

Users can perform different roles at different clinical sites within the same study, and can perform different roles for different subjects within the same site, but a user cannot perform different roles for the same subject within the same study.

NOTE: This rule has been implemented because performing different roles for the same subject could lead to permission conflicts and confusion in interpreting the audit trail.

A user who performs different roles at different sites within a study, will be limited to the database permissions associated with the role they perform at each site. For example, a user might be able to enter medical history data for all subjects at one site but only view it for all subjects at a different site. This role differentiation only affects database permissions (i.e. permission to view, create, modify and delete data and metadata). Users with multiple roles have access to the superset of the data collection tool views, **DFdiscover** applications, **DFdiscover** reports and Study reports, derived from all of the roles they are assigned in one study. Thus these permissions can vary among studies but not within a study.

Permissions for **DFdiscover** studies are specified in the **Permissions** view.

Username	Password Expiry	2FA Status	Full Name	Affiliation	Admin
data_coordinator	2024-05-23		data_coordinator	DFnet	Admin
data_manager	2024-05-22		data_manager	DFnet	254
datafax	2024-05-22		DFdiscover Admin		DFdiscover
demo_user	2024-03-12		Demo User 1	DF/Net Research Inc.	
site_invest	2024-02-22	Required	site_investigator	DFnet	
site_monitor	2024-05-22		site_monitor	DFnet	

New / Find

Duplicate

Delete

Save

Revert

All Users

Filter by Roles

Wildcard Search

Add Bulk Users

Username: demo_user Status: Active New Password



User Profile Preferences **Permissions** History

Study or DFdiscover Administrator: None

Allow access to Image Router

Study	Sites	Subjects	Role Name
<input checked="" type="checkbox"/> 101 Acme Pharma Development	ALL	ALL	Unrestricted
<input checked="" type="checkbox"/> 152 Demo A Production	250, 300	ALL	Unrestricted
<input checked="" type="checkbox"/> 200 CDASH Development	ALL	ALL	Unrestricted
<input checked="" type="checkbox"/> 252 Demo A Development	250, 350, 450, 999	ALL	Unrestricted
<input checked="" type="checkbox"/> 254 Acceptance Test Study	ALL	ALL	data_coordinator

User permission specifications include the following elements:

<p>Study or DFdiscover Administrator</p>	<p>Only a select subset of users should be given this global privilege. Any user can be assigned DFdiscover or study administrative privileges. Click  the button located at the right end of this widget.</p>  <p>Choose the correct option (None removes any previous admin permissions) and mark studies to include as needed.</p>
<p>Allow access to Image Router</p>	<p>The Image Router displays document pages that could not be identified, and is used to route any CRF pages among them to the appropriate study database. Selecting this option gives the user permission to use this application.</p> <p>It is important that only coordinating site staff with an understanding of all studies have this permission.</p>
<p>Status</p>	<p>Mark the checkbox at the left end of each row in the permission table to activate the permissions specified by the row when checked, and deactivate them when the box is not checked. This can be used to temporarily revoke permissions without deleting them from the table.</p>
<p>Study</p>	<p>Enter the study number or select the study from the study list. If a valid study number is entered, the field automatically displays the study name. The study must be specified before the other cells of the table are entered.</p>
<p>Sites</p>	<p>The sites, at which the user performs the role, are specified by entering ALL or * (indicating all sites), or a range of valid study sites. e.g. 1,5-8,98,99</p>
<p>Subjects</p>	<p>The subjects, for which the user performs the role, are specified by entering ALL or * (indicating all subjects at the specified sites), or a range of valid subject IDs . e.g. 1001-1099,1301-1399</p>
<p>Role Name</p>	<p>Type the study role name or pick it from the list of study roles. Since roles are defined separately for each study, the study must be specified before a valid role can be entered.</p>

A row highlighted in red indicates one of the following errors:


- An invalid entry in at least one of the cells in the row.
- Conflicting study roles, indicating that different roles have been specified for some set of study subjects within one or more sites.

Adding Study Permissions


New study permissions are added as follows:

1. Select the first empty row in the permissions table.
2. Specify the **Study**.
3. Specify the **Sites**.
4. Specify the **Subjects**.
5. Specify the **Role**.
6. If none of the values appear in red the row will be accepted as valid, the delete (trash can) icon will appear at the right end of the row, and a new empty row will be added to the permissions table.
7. When all of the rows needed to specify the user's study permissions have been completed, select the **DFadmin** menu item **File** > **Save** to save the current permissions. Changes made to user permissions are not saved until they are saved explicitly using **File** > **Save**.

Deleting Study Permissions

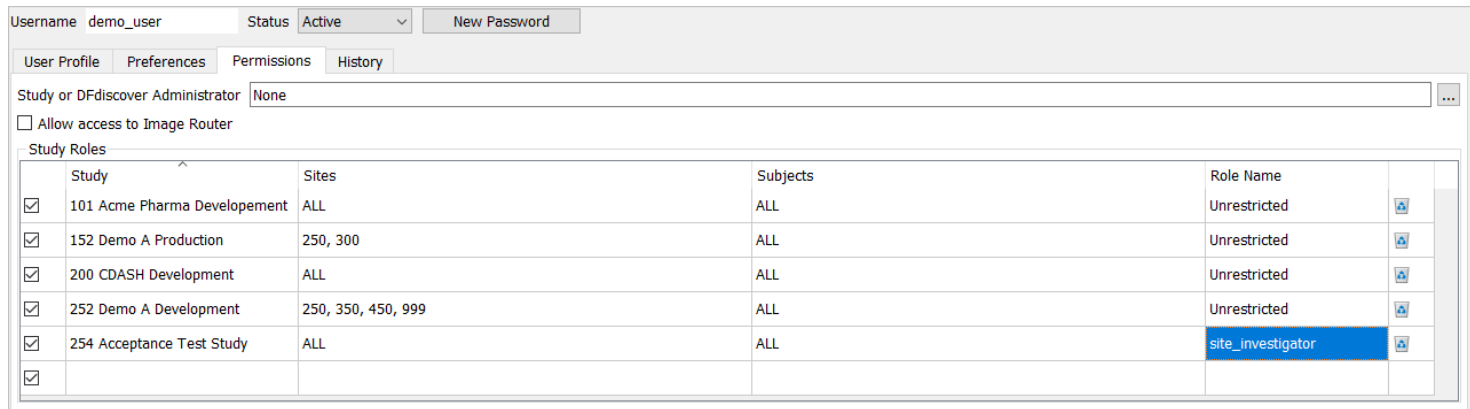
To delete a specific database permission specification (table row), click the corresponding delete icon (). A confirmation dialog is displayed, warning you that the user permissions are affected by this modification to the role definition. Click **Delete** to confirm and proceed with the deletion; otherwise, click **No** to cancel

the deletion.



Modifying Study Permissions

Study permissions can be changed by clicking and overwriting a cell and then pressing **Enter** or **Tab** to verify that the value is valid. Rows highlighted in red indicate errors.



If a set of permissions containing errors is saved the rows containing errors are dropped if they were newly defined, or revert to their previous valid values if the error was introduced by a modification.

ePRO Admin Permissions

Additional permissions may be defined per user for the management of a study's ePRO users, which is performed in **DFweb** on the ePRO Management page. These permissions are specified in the **ePRO Admin Permissions** view for each user.

There are three types of ePRO management permissions that can be assigned to individual users:

1. **Admin**: Full access to register participants, send invitations and login links, update participant details, and view ePRO user change and access logs.
2. **Monitor**: View access to the ePRO management page for monitoring, including access to personally identifying information (PII) such as participant name and email address.
3. **Viewer**: View access to the ePRO management page for oversight, where PII including participant name and email address are not displayed.

Username	Password Expiry	2FA Status	SSO Only	Full Name	Affiliation	Admin
data_coordinator	2025-10-27			data_coordinator	DFnet	
data_manager	2026-04-14			data_manager	DFnet	254
datafax	2026-02-17			DFdiscover Admin	DFdiscover	
datafax_atk	2022-12-04			ATK test account	DFdiscover	
demo	2026-03-18			Demo Account	DFnet	DFdiscover
demo_user1	2025-05-04			Demo User	DF/Net Software ULC	101,252

New / Find
 Duplicate
 Delete
 Save
 Revert
 All Users
 Filter by Roles
 Wildcard Search
 Add Bulk Users

Username: Status: New Password:

User Profile | Preferences | Permissions | ePRO Admin Permissions | History

Study	Type	Sites
<input checked="" type="checkbox"/> 804 DFdiscover 2024 Demo	Viewer	903-905
<input checked="" type="checkbox"/> 804 DFdiscover 2024 Demo	Monitor	901
<input checked="" type="checkbox"/> 804 DFdiscover 2024 Demo	Admin	902

ePRO Admin permission specifications include the following elements:

Status	Mark the checkbox at the left end of each row in the permission table to activate the permissions specified by the row when checked, and deactivate them when the box is not checked. This can be used to temporarily revoke permissions without deleting them from the table.
Study	Select the study from the study list. The study must be specified before the other cells of the table are entered.
Type	Select the access type from the list. The available types are Admin (full access), Monitor (view access with PII), Viewer (view access, no PII).
Sites	The site(s) for which the user has the access type defined are specified by entering one or more valid study site numbers, e.g. 1,5-8,98,99

A row highlighted in red indicates one of the following errors:

- An invalid entry in at least one of the cells in the row.
- Conflicting access, indicating that different access types have been specified for the same one or more sites within a study.

User History

DFAdmin [demo@explore.dfdiscover.com]

File Edit Action Help

Studies Roles Users Status License Master Traffic

Users (15) All Administrators

Username	Password Expiry	2FA Status	Full Name	Affiliation	Admin
data_manager	2024-07-24		data_manager		254
datafax	2024-05-14		DFdiscover Admin	DFdiscover	
datafax_atk	2022-12-04		ATK test account	DFdiscover	
demo_user1	2024-08-06		Demo User	DF/Net Software ULC	101,252
james530	2024-05-22	Required	James Brenton Moore	DFdiscover	
laura	2024-03-20		Laura Joldersma	DF/Net Software ULC	DFdiscover
macdouc_test	2023-11-13		Craig Macdougald	1,20,254	
pjansen	2024-06-09		Peter Jansen	DFnet	DFdiscover
samer	2024-07-17		Samer Chaudhary	DF/Net Software ULC	666,669,803,806
site_invest	2024-07-24		site_investigator		254
site_monitor	2024-07-24		site_monitor	DF/Net Research Inc.	254

Username: Status: New Password:

User Profile | Preferences | Permissions | History

Simplified Detailed history Refresh Save As...

Date and Time	Detail
1 2017-11-16 09:44:00	datafax added user "datafax" with Status "Active", Full Name "DataFax Admin", Language "English", Messages Preference "Mail", Router Access "No", Admin Access "No"
2	datafax assigned user "datafax" one-time password with Expiry "2017-11-16"
3 2017-11-16 13:22:01	datafax changed password with Expiry "2018-02-14"
4 2017-11-20 13:23:50	datafax changed user "datafax" with Status "Active", Language to "English", Messages Preference to "Mail", Router Access to "No"
5	datafax granted user "datafax" access to study "252 - Demo A Development" role "unrestricted" with Status "Active", Sites "All", Subjects "All"
6 2017-12-04 10:25:55	datafax granted user "datafax" access to study "152 - Demo A Production" role "unrestricted" with Status "Active", Sites "All", Subjects "All"
7 2017-12-15 16:49:42	datafax changed user "datafax" with Status "Active", Affiliation to "DF/Net", Street to "55 Head St, Suite 403", City to "Dundas", State to "Ontario", Country to "Canada", Telephone to "905 522 3282 x227"
8 2018-01-06 17:40:54	datafax changed user "datafax" with Status "Active", Postal Code to "L9H3H8"

The History tab displays the history of changes to the selected user, including user creation and modification, study access assignments, and one-time password assignments and password changes. The history displays the date and time of the change, the username of the person making the change, and the details of what was changed. History can be exported to an external file in Excel, PDF, HTML, or CSV format.

NOTE: When upgrading to **DFdiscover** Version 5.8.0 or later, existing user and role history must be loaded into the new admin history database for each study, using **DFadmindb** (see [Programmer Guide, DFadmindb](#) for details). Until the **DFadmindb** utility is run, it is not possible to view User, Role, or Study History in **DFAdmin**.

The default view is **Simplified**, where changes are grouped by date and time and described in a single statement. Select **Detailed history** to view the changes outlined in multiple rows and columns according to the specifics of the changes made. The detailed view is designed for easier use in Excel (filtering and sorting).

Changes are sorted from earliest to latest. Click the **Date and Time** column header in the table to reverse the sort order.

Additional actions available in the History tab include the following:

- **Refresh**: See the most recent changes since the table was loaded.
- **Save As...**: Export the current History view to Excel, PDF, HTML, or CSV format.

User profile details (name, email, etc.) for ePRO users are not displayed in the user history in **DFAdmin** to protect personally identifying information (PII) of study participants. The history of changes to ePRO user profile data is available from the **DFweb** ePRO Management page.

User and Session Settings

The **DFdiscover** client applications preserve settings for each user so that they can be retrieved and applied during subsequent login sessions. These settings include previous **DFdiscover** servers, proxy settings, main window sizes and positions, and the settings in application preference dialogs. User passwords are never included.

The settings are stored on the local computer in a unique location defined by the user login account. The settings are also separated by application and, where appropriate, by study.

Specifically, the settings are stored locally in these locations:

- **macOS and Linux**: `$HOME/.config/$TOOLNAME_db/$TOOLNAME.db`, where `$HOME` is the login directory for the user.
- **Windows**: `%APPDATA%\$TOOLNAME_db\$TOOLNAME.db`, where `%APPDATA%` is generally the `*AppData` sub-folder of the user's login directory.

`$TOOLNAME` is one of *DFexplore*, *DFsend*, *DFATK*, *DFAdmin* or *DFsetup*.

Generally speaking, users do not know, or need to know, that these settings files exist.

Sharing Settings File(s)

It is possible to streamline or standardize the settings for multiple users by creating one or more session setting files and then sharing those files with (new) users. This will give them the advantage of having certain settings pre-configured and standardized.

It is possible to share the settings file between the three supported operating systems. So it is possible to define the settings with an application on macOS and share that same file with users on **macOS, Windows and Linux**.

The file to be shared must match the `$TOOLNAME`. It is not possible to create a settings file for one application and share it for another application - this will certainly cause problems and/or unexpected behavior.

Local Cache

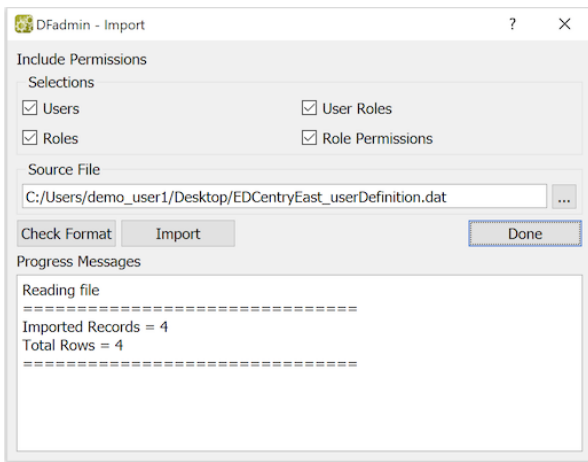
If the **DFdiscover** or study administrator has enabled local caching for one or more studies, frequently requested properties such as study logos and configuration files will be automatically cached in each user's local environment. Any user can clear their cache by selecting **File** > **Clear Cache** within **DFexplore** or by removing the following file from their local filesystem. These two actions are equivalent. Under normal circumstances there is no need for the user to clear or remove the local cache - **DFexplore** quietly manages the cache without user intervention.

NOTE: Study data is *never* cached.

- **macOS and Linux**: `$HOME/.config/DFlocal_db/$DFlocal.db`, where `$HOME` is the login directory for the user.
- **Windows**: `%APPDATA%\DFlocal_db\DFlocal.db*`, where `%APPDATA%` is generally the `*AppData` sub-folder of the user's login directory.

Import

Study roles, role permissions, role restrictions, role assignments to users, and user contact information can be imported from an ASCII file. Import will overwrite any matching specifications that already exist. The import dialog includes options for selecting the type of records to be imported (i.e. Users, User Roles, User ePRO Permissions, Roles and Role Permissions and Restrictions) from the import file as illustrated below. When **Role Permissions** is selected, both Role Permissions and Role Restrictions are imported. The import file must be formatted as described in [DFuserdb.log](#), excluding the Record Time Stamp and Modifier fields.



Importing Users, Roles and Permissions

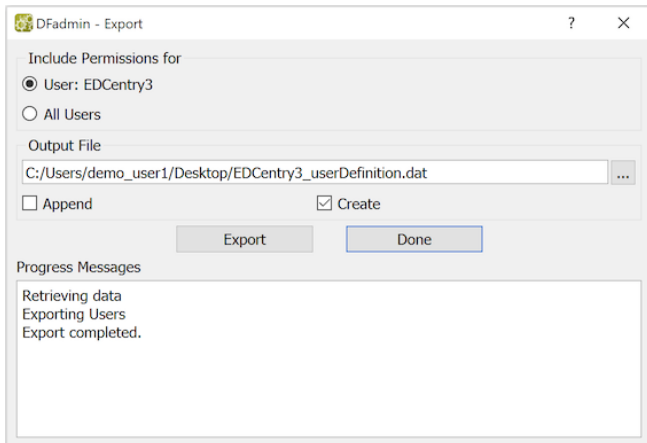
1. Select **File** > **Import**
2. In the **Selections** section of the dialog select the types of data you wish to import.
3. Specify the **Source File** name or pick a file using the file selection button.
4. Click **Check Format** Any invalid entries will be displayed. **Import** will be enabled only if there are no invalid records.
5. Click **Import**. Progress and errors messages are displayed while importing.
6. Click **Done** to close the Import dialog.

Export

All user account information, including study permissions, can be exported to an ASCII file, with the output format described in DFuserdb.log.

Exporting Users

1. Optionally, highlight the user to be exported (if export is to be limited to a single user).
2. Select **File** > **Export**.
3. In the dialog, select the current user or all users.
4. Enter an **Output File** name or pick the file using the file selection button.
5. Choose **Append** to write the exported data to the end of file, or **Create** to overwrite the file with the exported data.
6. Click **Export**. Progress and errors messages are displayed while exporting.



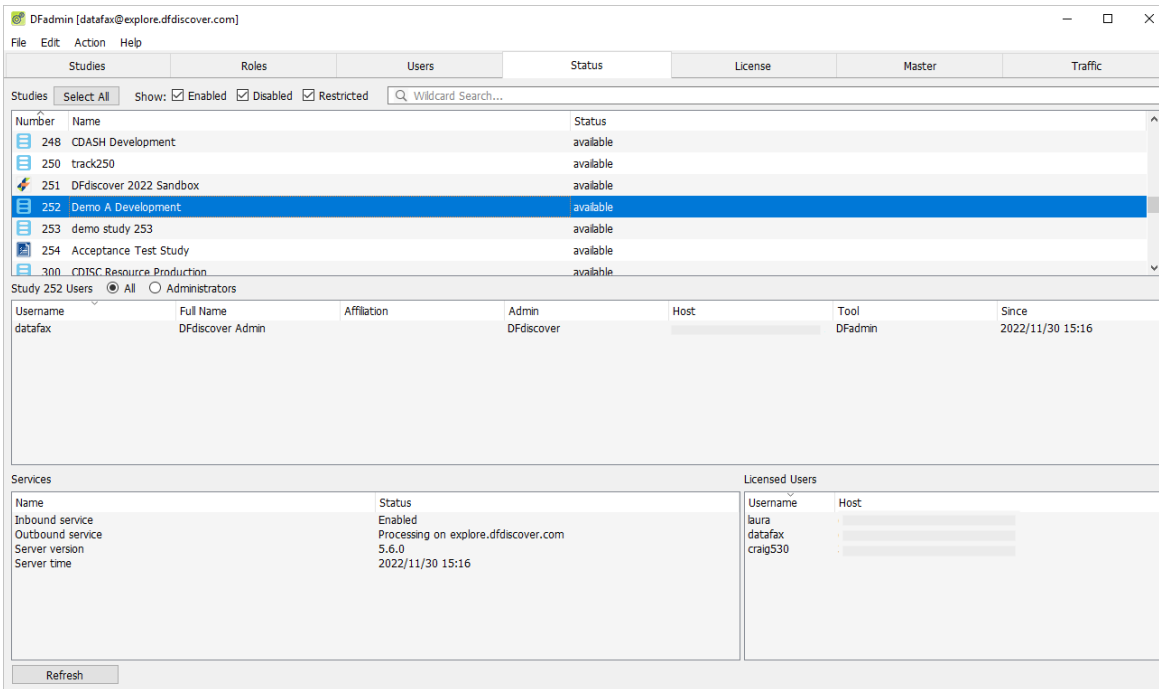
7. Click **Done** to close the Export dialog.

Status

The **Status** tab shows the studies, users, and services currently running on the **DFdiscover** server. This includes:

- the number, name and status of all studies registered on the **DFdiscover** server which the user has permission to administer


- this list may be restricted to any combination of: enabled, disabled and restricted studies using the check boxes above the list; or use **Wildcard Search** to find studies by their number or name
- the current status of each study is indicated with an icon: enabled studies (📄), disabled studies (🚫), and restricted studies (🔒).
- for disabled and restricted studies the user who set the status and the reason message they entered equal to or less than 150 characters is also displayed
- if a study is in read-only mode **Study is in Read-Only mode** is displayed in red when the study is selected
- when a study is selected **All** current users or just **Administrators** are shown in the middle panel
- the current status of the Inbound and Outbound services appears in the bottom left panel
- all users currently using a **DFdiscover** license (seat) across all studies and services appear in the bottom right panel



Changing Study Status

Access to each study can be limited using the Action menu to:

- Disable a study, which stops the study database server making it unavailable to all users, essentially locking the study database so that neither the data nor the study setup can be modified.
- Restrict a study, which makes it available to **DFdiscover** and study administrators but unavailable to all other users.
- Make a study Read-Only, which allows study data and setup information to be viewed but not changed.



A study may be simultaneously **Restricted and Read-Only** which limits study access to **DFdiscover** and study administrators while preventing them from making any changes. This is accomplished by first 'Restricting' the study and then making it **Read-Only** or vice versa, the order in which these 2 operations are performed is irrelevant.

Putting a study into **Restricted** access mode does not take effect immediately. Existing users can continue working until they log off. It does however prevent new users from connecting to the study. When all existing connections have terminated, the administrator who put the study into restricted mode receives an email message indicating that the study is now restricted to **DFdiscover** and study administrators.

It is not possible to **Disable** a study while it is being used. In such cases the study should first be restricted and then disabled when all users have logged out.

It may take several seconds for the **Disable** action to complete, as the study server may need to perform extensive shutdown processing before reporting that

the action has been successful.

Selecting **Enable** returns studies which have been disabled or restricted to full access, and selecting 'Read-Write' returns studies which have been made **Read-Only** to full access.

It is possible to change the status of more than one study at a time by selecting multiple studies and then selecting the desired status from the **Action** menu. Use **Shift**-click to select a range of studies, and **Control**-click(Win) or **Command**-click(Mac) to select studies one at a time.

When a change in status is selected a confirmation dialog will appear with **OK** to proceed and **Cancel** to abort the operation.

When a study is **Disabled** or **Restricted** the confirmation dialog allows entry of an optional message which if provided will be displayed in the study selection dialog in **DFexplore** and **DFsetup**. This message is also displayed in the **Enable study** confirmation dialog to remind users why the study is in it's current state before they reverse it. The reason for disabling or enabling a service have been standardized to accept a maximum 150 character message.

Study status cannot be changed if anyone is currently using the study in either **DFexplore** or **DFsetup**. This applies not only to status changes that limit access, but also to those that reverse previously applied limitations.

NOTE: When a study is disabled the study database and setup remain unchanged, but client applications, and image processing daemons will not be able to communicate with the study server.

Images that arrive while a study server is disabled will be received and processed by **DFdiscover** but any pages barcoded for a disabled study will be held until the study is re-enabled, and will not show up in the study's new image queue until processing is triggered by the arrival of the next image received by the **DFdiscover** server. It is not necessary for the next image to contain pages barcoded coded for the study in question or for any study.

When a Query Report is successfully sent to its destination clinical site, the study database server is notified. This allows it to update the status field in all queries included in that report, to indicate that they have been transmitted. Since updating the Query database requires the study server, queries will not be updated if the study server is disabled while Query Reports are being transmitted.

NOTE: While a study is in restricted access mode, **DFdiscover** and study administrators will be able to use client and shell-level applications to the full extent allowed by their study role permissions, while other users will not be allowed to use any of the **DFdiscover** applications, including: **DFexplore**, **DFsetup**, **DFexport.rpc**, **DFimport.rpc**, **DFprintdb**, **DFcompiler** and **DFstatus**. Also, since shell-level applications like **DFexport.rpc** often appear in study specific reports and custom applications, this restriction will extend to applications that may normally be able to run from a UNIX login account.

Restricting a study has no effect on the processing of incoming or outgoing images, except that only a **DFdiscover** or study administrator will be able to save records in **DFexplore** Image View and route unidentified pages to the study in the image router.

Unlike 'read-only' and 'disabled', 'restricted' access is a volatile state. If the **DFdiscover** master is shutdown and restarted, disabled studies remain disabled, and read-only studies remain read-only, but restricted studies become unrestricted.

NOTE: While a study is in Read-Only mode client applications and shell-level applications are available to view and export study data and setup information, but the study setup and database cannot be changed. Other components such as DFsas job files, custom list views, and data tasks cannot be created or changed.

Studies are typically put into Read-Only mode when data collection is complete, all queries have been resolved, the data has been verified against source documents, and database lock has been completed. Statistical analysis may be ongoing or an FDA audit may be in progress thus the study cannot be disabled, but no further changes are expected or allowed.

In this state we do not expect to receive any additional CRFs by email, PDF, or fax submission. If any pages barcoded with the **DFdiscover** study number are received they are sent to the unidentified image router where they can be reviewed, but they cannot be routed to the study database as long as it remains in read-only mode.

Traffic

Traffic is the general description for three related features:

- incoming document processing, known as Inbound
- outgoing document processing, known as Outbound
- received document receipt processing, known as Auto Reply

These features are all administered in the Traffic tab of **DFadmin**.

DFdiscover can be configured to send a receipt back to selected clinical sites each time that a document is received from them. Also, incoming document processing can be blocked from any sender that is not registered with **DFdiscover**. Both features are administered in **DFdiscover** through an option known as **Auto Reply**.

Auto Reply is intended as a debugging aid within **DFdiscover**, assisting in the diagnosis of clinical sites from whom document reception is unreliable. Once the problem is diagnosed it is no longer necessary for this feature to remain enabled for that site.

Incoming Documents

There is a great deal of work that **DFdiscover** does to automatically process and route incoming documents. In this section, the process is described in detail.

Incoming and outgoing faxing from **DFdiscover** is accomplished by third-party software, **HylaFAX**. For purposes of this description, both are generically referred to as the "Fax software". The **DFdiscover** Inbound and Outbound services act as overseers and provide additional end-user functionality and rely on the Fax software to interact with the modem devices.

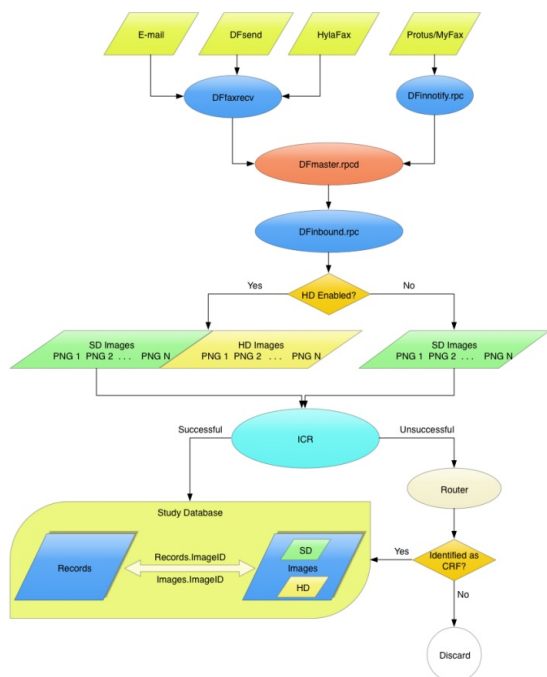
IMPORTANT: Faxes and Email

In addition to receiving faxes via the traditional physical fax machine, **DFdiscover** is also capable of receiving incoming documents via email, **DFsend**, and **DFexplore**'s Submit PDF function.

Once a fax or document has been received to disk, the handling of both input streams is identical. This alternate input stream is described in [Processing of an incoming email](#).

DFdiscover relies on an email message from the **HylaFAX** software to initiate its incoming fax processing. Incoming fax processing by **DFdiscover** is event driven; the processing software executes only when a fax has arrived. If there are no recently arrived faxes requiring processing, no Inbound service will be running. Each time a new fax arrives, the Inbound service is dispatched by the master daemon to process that fax. If multiple new faxes arrive at or near the same time, the master daemon will dispatch the first fax to the Inbound service, queue the additional faxes, and then de-queue each fax as the Inbound service completes the previous fax. Incoming faxes will never be discarded because there is no Inbound service available. The fax will be queued by **DFdiscover** (freeing **HylaFAX** to receive the next fax) until the service is available.

DFdiscover routes an incoming document from the Fax software, email, or **DFsend** to its ultimate destination within **DFdiscover**. The remainder of this section describes each step in detail.



Processing of an incoming document

1. Document arrives and is received by the Fax software

- o **HylaFAX** configurations: When a completed CRF containing *N* pages is sent from a fax machine at an investigator site to a fax modem running the Fax software, the fax modem answers the call and creates a file for receipt of the fax document in the **DFdiscover** incoming fax directory. As receipt of the fax progresses, the size of the incoming file increases. Transmitted in standard resolution, a typical fax page requires approximately 25 Kilobytes of disk space.

When transmission is complete, a mail notification message is sent to user *datafax*. Receipt and processing of the fax transmission by the Fax software is complete.

If the transmission enters the system via an email message, alternate processing takes place during this step and the next. That processing is described in [Processing of an incoming email](#).

The email message sent by the Fax software (or received directly via email) is parsed by the **DFfaxrecv** application that checks the message content for validity. The message content is valid if the following conditions are true:

1. the sender of the message is the Fax software
2. the subject of the message indicates a newly arrived fax

If the message content is valid, the filename is then passed to the master to signal it that this file needs processing.

2. DFdiscover master assigns the Inbound service to process the document

In dealing with incoming documents, the responsibility of the master daemon is to assign the document to the Inbound service.

If the Inbound service is busy processing another document, the master records the incoming document's name in a FIFO queue for future processing.

Subsequently, when the Inbound service completes processing of the previous document, the master dispatches the first document from the queue.

As soon as the Inbound service starts, it sends the master a message to indicate that it has started to process the assigned document. The master updates its inbound database to indicate that the service is processing.

3. DFdiscover identifier is assigned to the document

Immediately after the Inbound service has notified the master that it has started processing the assigned document, it requests a unique **DFdiscover** identifier for that document from the master. The master consults sequence files in `/opt/dfdiscover/work` to determine the identifier.

From the document name, the master determines the year and week in which the document arrived. It uses this to construct a 4-digit number, *YYWW* where *YY* is the last two digits of the year (minus the century) and *WW* is the week within that year (01 to 53).

The master searches for a file named `.seqYYWW` in `/opt/dfdiscover/work`. If no such file exists, the master returns the identifier of 1 to the incoming daemon and creates the sequence file `.seqYYWW` recording 2 as the file contents. If the file does exist, its contents are returned as the **DFdiscover** identifier, and then incremented by one before being written back to the file.

The **DFdiscover** identifier supplied by the master to the Inbound service represents the sequencing of when a document arrived within a particular year and week. For example, the third document to arrive during the week of September 16, 2017 would have a *YYWW* value of *1738* (the 38th week of 2017) and a sequence number of *0003*. The unique **DFdiscover** identifier for that document then becomes *1738/0003*. Because of the algorithm used to generate the identifier, **DFdiscover** can guarantee that the identifier is unique within the **DFdiscover** installation.

For the remainder of this description, *YYWW/FFFF* will be used as the example **DFdiscover** identifier.

Another point to note here is that these unique document identifiers are being assigned before the document is processed, and thus before there is any possibility of determining which study the document belongs to. As a result all document identifiers are unique, not just within each **DFdiscover** study, but across all **DFdiscover** studies being managed on the same **DFdiscover** installation.

IMPORTANT: It is extremely important that the sequence files, used to assign these unique document identifiers, not be manually edited or deleted. Deleting a sequence file will cause the sequencing for that *YYWW* combination to be reset to 1, which will result in non-unique document names. Fortunately, **DFdiscover** will detect this condition and prevent old document from being overwritten by new document with the same name. Even after a particular *YYWW* has passed in time, it is important to leave the sequence files in case it is ever necessary to retrieve and re-process a document from the archives.

4. Inbound service archives the document

If the value of `INBOUND_ARCHIVE_DURATION` in the daemon's configuration is any number greater than 0, a copy of the received document is made (without alteration to the document) in the archive directory. The location of the archive directory is determined by evaluating the `INBOUND_ARCHIVE_DIR` parameter in the daemon's configuration. The name assigned to the archive copy is the *YYWW/FFFF* identifier.

5. Inbound service decompresses the document

The Inbound service breaks the original *N*-page document into its individual *N* pages. This occurs in the incoming daemon's work directory.

The conversion creates 100 dpi PNG files, as well as 300 dpi PNG files if High Definition (HD) images are enabled in **DFadmin**. There is exactly one PNG file created for each page in the document. Hence an *N*-page document creates *N* 1-page PNG files. The page names for the PNG files are determined from the **DFdiscover** identifier. In the Inbound service's working directory, the sub-directory *YYWW* is created (if necessary) containing files named

```
0.FFFF1 0.FFFF2 0.FFFF3 ... 0.FFFFN
```

The files are then renamed so that the page number within the document is leading zero-padded to three digits.

IMPORTANT: An important point to be aware of here is that the original documents cannot be archived in a separate directory for each study. The main reason for this is that there is no guarantee that **DFdiscover** will be able to uniquely link a single document with a single study. A site involved in multiple studies may send CRFs for more than 1 study in the same transmission. You might even use more than one **DFdiscover** study number to keep separate databases (e.g. lab data, clinic data, endpoint adjudication reviews, etc.) for the same study. Also a document might not contain any barcoded CRFs. Consequently all received documents across all studies are archived to the same document archive directory.

Conversion of the original document to PNG files is complete. The Inbound service indicates this to the master by sending it a message containing the name of the archived document, the number of pages it contains, the document's receipt date and time stamp (or the email transmission's receipt date and time stamp), the sender identifier of the document (for faxes, this is transmitted by the sending fax machine; for emails, this is the sender's email address) and whether or not the document was archived. The master records this information by appending it to the file `/opt/dfdiscover/work/fax_log`. This log file is in plain text format and can easily be consulted to determine a particular document's context. In fact, this is the file that is consulted by client applications like **DFexplore** to determine context.

6. Each page is put in standard DFdiscover orientation

At this point in the incoming processing there are one or more PNG files named *0.FFFFPPP*. Processing now proceeds sequentially and repeats separately for each page. If page processing fails at any step, the PNG file is moved to the unidentified image routing directory, `/opt/dfdiscover/identify`, where the page can be subsequently viewed and dealt with manually using the **DFexplore** Image Router function.

The width of the page is tested to ascertain that it is within a reasonable tolerance of being a US-letter or A4 page in portrait orientation. Then the height of the page is tested to be within the same tolerance.

The service then begins searching for the **DFdiscover** barcode signature line near the top of the page (within the top 2.5" of the page). If the line cannot be found, it is assumed that the page is wrong-side up, it is flipped right-side-up and the line search is repeated.

After the signature line is found, the vertical offset between the right end and left ends of the line is used as the measure of skew inherent in the entire page. An angle of rotation is calculated using the arctangent of the vertical offset divided by the length of the line. The entire page image is then rotated by

this angle around the left corner of the signature line. The assumption here is that the same degree of correction is required across both the length and width of the page. Occasionally, this assumption fails, and skew correction is not as good as it might be, but in general the correction is adequate and better than if no correction had been applied at all.

When rotation completes, the resulting page image is then shifted so that the top of the signature line is coincident with the top of the page and the left end of the signature line is 0.5" from the left edge of the image. This is the standard **DFdiscover** orientation for a CRF image. All images must be in this orientation before they can be scanned for their **DFdiscover** barcode.

7. Inbound service scans page for barcode

In standard **DFdiscover** orientation, the barcode for a CRF can be found starting an inch from the left edge of the image and down 0.125" from the top of the image. A readable barcode contains either 17 or 26 vertical bars, evenly spaced at 0.25" intervals, each of thickness 0.04" or 0.12". The 0.04" bars are counted as having value 1 while the 0.12" bars have value 0. The concatenation of 0s and 1s for the first 8 bars is evaluated as a binary number and assigned to the study number, the next 9 bars constitute the plate number, and, if present, the last 9 bars constitute the sequence number. For more information on **DFdiscover** CRF barcoding, consult [Barcode User Guide](#).

The newly rotated and shifted page image is then written to a file named *FFFFPPP* and the original version of the page in *0.FFFFFPP* is deleted. At the same time, a placeholder record for the data that will be ICRed from the image is written to the file *DFrecin.###* where *###* is the study number that was scanned from the initial part of the barcode. If *DFrecin.###* already exists, the new record is appended to the file; otherwise, the file is created. Although unlikely, it is possible for one document to contain CRFs from multiple studies and hence multiple *DFrecin.###* files may be created. The purpose of each *DFrecin.###* file is to temporarily record the:

- document page name
- study number
- and plate number

in preparation for the next phase, ICR, of the Inbound service.

As previously noted, pages which fail any of the preceding steps are moved to the unidentified images routing directory, */opt/dfdiscover/identify*. Pages that have passed all of the steps will be ICRed and then added to the new pages input queue for a study, where they can be retrieved and reviewed in **DFexplore**, and finally entered into the study database.

8. ICR processing occurs for each identified page

When all pages have completed the barcode reading step, the ICR software is launched by the incoming daemon. ICR processing begins by scanning for file names that begin with the *DFrecin* prefix. For each such file that is found, the study number is determined by stripping the last three digits from the file name. The Inbound service then attempts to make a connection with the database server for that study number. If the connection is successful, the service requests the location of the ICR template file and the CRF pages directory from the study database server.

For each record in the *DFrecin* file, the plate number is extracted from the 5th field and the image identifier is extracted from the 3rd field. These two pieces of information together with the already read ICR template permit the software to scan the CRF image for data fields.

ICR can read hand-written numbers, check boxes, dates (containing digits and characters), and visual analog scales, but not hand-written text. ICR will fail to read a value and leave the field's value as blank, if any one or more of the following conditions is true:

- The field type is string, meaning that the contents are hand-written text
- The field cannot be located on the CRF
- The read value is illegal when tested against the field's legal range
- More than one choice has been marked for a choice field

IMPORTANT: ICR is not data entry. The ICR step should be viewed only as an assist to data entry, not as a complete replacement for it. It is important that all records be reviewed to complete blank fields, correct any ICR errors, pick up corrections and comments written by investigators in the margins of the CRF, and flag missing data with queries. All of these tasks can be performed by study personnel through **DFexplore**.

9. Pages and data are sent to the study database

When ICR processing completes on a CRF page, the resulting data record is assigned a workflow level of 0 and is sent via a network message to the study database server, which appends it to the study input queue for new images file. The corresponding image file is moved from the Inbound service's working directory to the study's CRF pages directory. The corresponding record from *DFrecin* is then deleted. If a connection to the study server cannot be established (i.e., the study server may be in read-only mode), the CRF images are sent to the **DFdiscover** *identify* directory, instead of to the study's *pages* directory.

When all of the records in a *DFrecin* file have been processed, the file will be empty and it is deleted. Since the image files have also been moved to the study's CRF pages directory, the Inbound service working directory is empty and is deleted. The service will report an error if this condition is not true.

Incoming document processing by **DFinbound.rpc** is now complete. Before termination, the Inbound service sends a completion notification message to the master and then exits.

Processing of an incoming email

1. Document arrives by email

DFdiscover is also capable of accepting CRFs that are received by email. The following requirements for emailed CRFs must be met:

- The recipient of the message must be user *datafax* on the **DFdiscover** server. If emails are arriving from external locations, email routing must allow email to reach the **DFdiscover** server.
- The sender of the message is not user *datafax* or any email address that contains the phrase *MAILER-DAEMON*.

- The body of the message must contain one or more PDF or TIFF files encoded as MIME attachments. Other attachment types will be silently ignored. If no attachments can be processed, the email will be bounced with a Data Format error.
- Each TIFF attachment may contain one or more CRF pages. The TIFF images may be color, grayscale and black-and-white, at least 100 dpi resolution, and no more than 300 dpi resolution.
- The mail application (typically **sendmail** on the UNIX platform) must have permissions to write to the `/opt/dfdiscover/incoming` directory. Typically this means that the directory must have access and write permissions for everyone, unless the user id (`root`, `uucp`, or `daemon`) that executes the mail application belongs to the `studies` group.

2. Message body is decoded

The decoding is performed by the **DFdiscover DFfaxrecv** application which breaks out the individual attachments. TIFF attachments are saved to `/opt/dfdiscover/incoming` with names in the form `email#####-###.tmp`. PDF attachments are saved to `/opt/dfdiscover/incoming` with names in the form `email#####-###.pdf`. These PDF files are then passed to **DFgs** to convert them to TIFF files with names in the form `email#####-###.tmp`.

3. DFdiscover master is notified

The **DFfaxrecv** application notifies the master of the arrival of the new images. If **DFdiscover** is not operational at the time, the transmission remains queued in this directory. It will subsequently be de-queued and processed automatically by **DFdiscover** when it restarts.

4. Regular processing resumes

The remainder of the processing of the email transmission follows the same procedure already described in [Processing of an incoming document](#).

Outgoing Documents

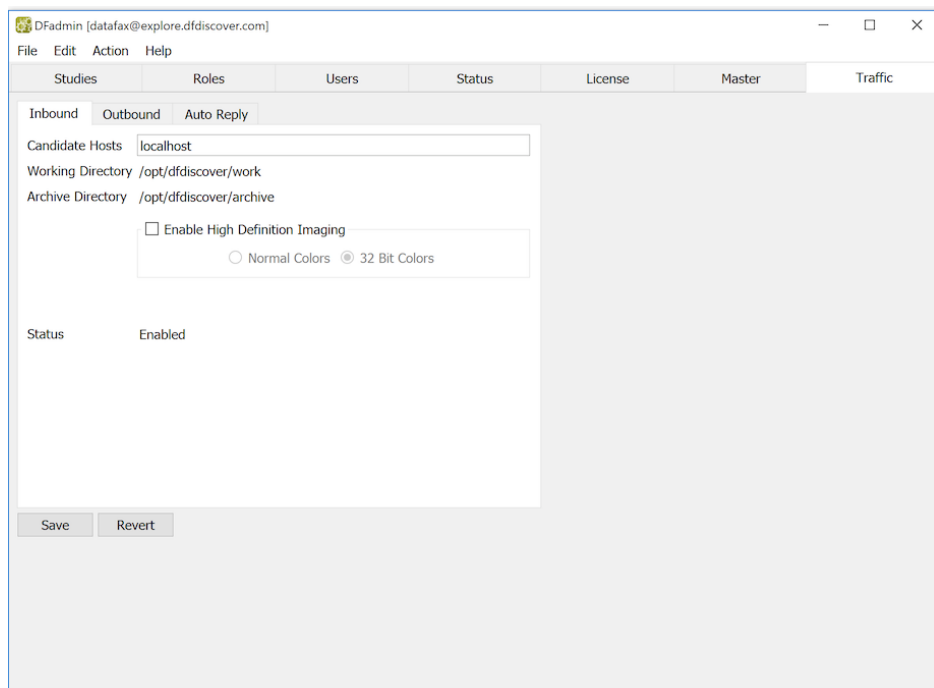
Through **DFdiscover**, it is possible to send documents to one or more recipients directly from your computer. Users will exercise this feature most via **DF_QCfax**, one of the standard **DFdiscover** reports, which sends Query Reports to the participating sites of a study.

DFdiscover manages outgoing traffic with the Outbound service that behaves like the UNIX **lpd** daemon. It is always running in the background (unlike the Inbound service which only executes when there is an incoming document to process), and it works from a spool directory. Programmers can interface with the outgoing daemons via the shell-level commands **DFsendfax** (to queue a document for sending), **DFfaxq** (to monitor the status of the outbound queue), and **DFfaxrm** (to remove a document from the queue). These **DFdiscover** utilities are described in [Programmer Guide, Shell Level Programs](#).

Inbound

Documents that arrive at one of the fax modems on a **DFdiscover** server are received by **HylaFAX**, which subsequently informs the **DFdiscover** master daemon that there is a document that needs to be processed. The master then assigns the document to the incoming daemon for processing. These are the basic actions of the Inbound service.

The Inbound feature must be defined before **DFdiscover** will be able to process incoming documents. The definition and configuration of the Inbound feature is done in the **Inbound** view.



Modifying Inbound Configuration

When **DFdiscover** is started for the first time, **DFdiscover** will identify that there is no Inbound service defined and it will define and configure it with default values. In most environments, these default values will not need to be changed. However, if re-configuration is needed, only the candidate host can be changed using **DFAdmin**:

- **Candidate Hosts:** Enter the host names (comma-delimited) of all computers that the master can consider as candidate CPUs where the Inbound service might be permitted to run. In 99% of all cases, only one hostname is needed and that is the hostname of the master.

IMPORTANT: **Candidate Hosts** also require slave daemons. The Inbound service can only run on machines that are also running slaves. If a candidate host is specified that is not also running a slave, that candidate will never be chosen.

- **Working Directory:** The working directory must be writable by user *datafax*, and is typically defined with the same value as the **DFdiscover** work directory, `/opt/dfdiscover/work`. To change its value, the inbound configuration file `/opt/dfdiscover/lib/DFinbound.cf` must be directly edited by the **DFdiscover** administrator.

The working directory is temporary space only - once an incoming document is processed, the working directory becomes empty again.

- **Archive Directory :** This is a directory where the Inbound service stores the original documents (TIFF or PDF). It must be writable by user *datafax*. Since the files may be large, this should be a directory with lots of available space. Also this directory should be monitored regularly to make sure that space is available for archiving of incoming transmissions (see [Archive File Maintenance](#) for more information). To change its value, the inbound configuration file `/opt/dfdiscover/lib/DFinbound.cf` must be directly edited by the **DFdiscover** administrator.

After making any changes, select **File** > **Save** to save the changes and re-configure the Inbound service. Modifications can be made regardless if the Inbound service is currently enabled. If the Inbound service is processing while its configuration is modified, the modifications will not take effect until the next incoming document is processed.

Enable/Disable Inbound


In the Inbound view, **Status** is a read-only field that shows the current status of the Inbound service. Typically, the service is always enabled, which means that it can process incoming documents as they arrive. However, the service can also be disabled by the administrator - when disabled, documents continue to arrive, but they wait in a queue and are not processed by **DFdiscover** until the Inbound service becomes enabled.

The possible statuses are:


- *Enabled*, on-line, and available for processing each incoming document
- *Disabled*, off-line, incoming documents are queued and not processed

To change the status of the Inbound service, select one of:

- **Action** > **Enable Inbound** Enable the currently disabled Inbound service.
- **Action** > **Disable Inbound** Disable the currently enabled Inbound service.

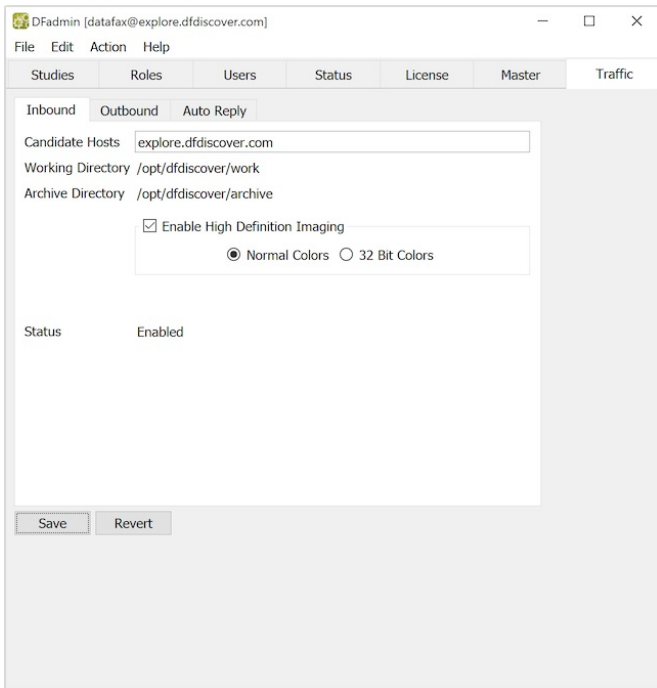


A reason can be provided to describe why the Inbound service is being disabled. A reason is not required but is recommended particularly if the server is administered by more than one person. When the Inbound service is disabled the reason appears in the **Status** field of the Inbound view.



Enable/Disable High Definition (HD) Imaging

DFdiscover accepts incoming documents via fax, email and **DFsend**. Documents that are transmitted via fax can have only one resolution, 100 dpi (SD, or standard definition) - this is part of the G3 fax standard. Documents that arrive via email or **DFsend** may have higher resolution if they were scanned at higher resolution at the sending site. For such documents it is possible to instruct **DFdiscover** to also image and store them on the server at a higher resolution, namely 300 dpi (HD, or high definition).



While it may seem preferable to image documents at the highest available setting, there are several considerations:

- High definition images require more disk storage space on the **DFdiscover** server.
- High definition images require more bandwidth, and hence time, to transmit to an **DFexplore** client. On the client side, each user may also choose whether to request HD or SD.
- **DFdiscover** will always keep an SD version of each page of each document. If high definition is also enabled, the HD version is also kept, thereby further increasing storage requirements.
- If the source documents are in color or contain shaded areas that need to be maintained, then high definition should be enabled. With HD enabled, one can also select the color fidelity in the imaged document. In general, the **Normal Colors** setting is acceptable for clinical study documents. Use **32 Bit Colors** only if the source document has photographic detail that needs to be preserved. If **32 Bit Colors** is enabled, the stored images are significantly larger (up to 10 times) than a SD image (and hence will also take longer to transmit) - be mindful of this before enabling. For most documents, there is no discernible improvement in quality with **32 Bit Colors** over **Normal Colors**.

In general, if the existing SD quality is sufficient, there is no need to enable HD.

When high definition is enabled or disabled, the new setting is applied only to documents that are received after the changed setting is saved. **DFdiscover** does not retrospectively re-image previously received documents.

There is a similar setting at the client end, in **DFexplore**. Those details are not described here but, in summary, the end user is able control whether they wish to request HD or SD images. Over low-speed or low-bandwidth connections, SD is more efficient and hence preferred. Selecting HD in such an environment can result in slower image loads and decreased **DFexplore** performance.

Outbound Service

Documents which require transmission from **DFdiscover** to external sites, typically participating investigative sites, are handled by the Outbound service.

The Outbound service must be defined and configured before **DFdiscover** is able to send outgoing documents. The definition and configuration of the Outbound feature is done in the **Outbound** view.

File Edit Action Help

Studies	Roles	Users	Status	License	Master	Traffic
Inbound	Outbound	Auto Reply				
Candidate Hosts <input type="text" value="explore.dfdiscover.com"/>						
Working Directory <input type="text" value="/opt/dfdiscover/outgoing"/>						
Polling Interval	<input type="text" value="60 seconds"/>					
Delivery Period	<input type="text" value="24 hours"/>					
Fax Service	<input type="text" value="EMAIL"/>					
	User ID <input type="text"/>					
	Password <input type="text"/>					
Status	Processing on explore.dfdiscover.com					
<input type="button" value="Save"/> <input type="button" value="Revert"/>						

Modifying Outbound Configuration

When **DFdiscover** is started for the first time, **DFdiscover** identifies that there is no Outbound service defined and it will define and configure one with default values. In most environments, these default values do not need to be changed. However, if re-configuration is needed, the following parameters can be changed:

- **Candidate Hosts:** Enter the host names (comma-delimited) of all computers that can be considered as candidate computers where the Outbound service might be permitted to run. In 99% of the cases, only one hostname is needed and that is the hostname of the master.

IMPORTANT: **Candidate Hosts** also require slave daemons. The Outbound service can only run on machines that are also running slaves. If a candidate host is specified that is not also running a slave, that candidate will never be chosen.

- **Working Directory:** The working directory must be writable by user *datafax* and is defined by default with the same value as the **DFdiscover** work directory, */opt/dfdiscover/work*. To change its value, the outbound configuration file */opt/dfdiscover/lib/DFoutbound.cf* must be directly edited by the **DFdiscover** administrator. It cannot be changed using **DFadmin**.

The working directory is temporary space only - once an outgoing document is sent, the working directory becomes empty again.

The permissions on the files queued for outgoing transmission are such that only user *datafax* can read or write them. This prevents any other user from being able to tamper with an outgoing document before it is delivered.

- **Polling Interval:** This value is the number of seconds that the Outbound service will wait between checks for new requests in the outgoing document queue. The Outbound service runs in a loop where it checks the queue for work, does whatever work may need to be done, then sleeps for the specified number of seconds before beginning the loop again. A recommended value is anywhere in the 20 to 60 second range.
- **Delivery Period:** This value is the maximum number of hours that the Outbound service will wait for a reply from **HylaFAX** as to the disposition of a document that has been queued for transmission. If a reply is not received in this time period, **DFdiscover** considers the transmission of that document to have been a failure; resources are freed, any failure commands for the document are executed, and the queue entry is removed. A reasonable value is in the 8 to 16-hour range.

If a large number of documents is typically being queued, be careful not to make this value so small that even under normal circumstances, **HylaFAX** would not be able to reply in time as it would still be processing the large volume of documents.

- **Fax Service:** You can select the Fax Service your **DFdiscover** server is setup for. **EMAIL** or **HYLAFAX** service can be used for document handling services. While available in the drop-down, **PROTUS** is no longer supported.
 - **EMAIL** Email services will be used for document handling. Ensure that external email is configured on your **DFdiscover** server - *postfix* is a common email service. Configuration of *postfix* or other email services is beyond the scope of this document.
 - **HYLAFAX** HylaFAX software will be used for fax handling. Ensure that HylaFAX is installed on the **DFdiscover** server.

IMPORTANT: **DFadmin** does not accept the `&`, `|` or `=` characters.

After making any changes, select **File** > **Save** to save the changes and re-configure the Outbound service. Modifications can be made regardless of whether or not the Outbound service is currently enabled. If the Outbound service is processing while its configuration is modified, the modifications will not take effect until the next outgoing document is processed.

Start/Stop Outbound

In the Outbound view, **Status** is a read-only field that shows the current status of the Outbound service. Typically, the service is always processing, which means that it can process outgoing requests as they are issued. However, the service can also be stopped by the administrator - when stopped, outgoing document requests are not processed until the Outbound service is re-started.

The possible statuses are:

- *Processing on hostname*, on-line, running on *hostname*, and available for processing each outgoing document request
- *Disabled*, off-line, outgoing document requests cannot be processed.

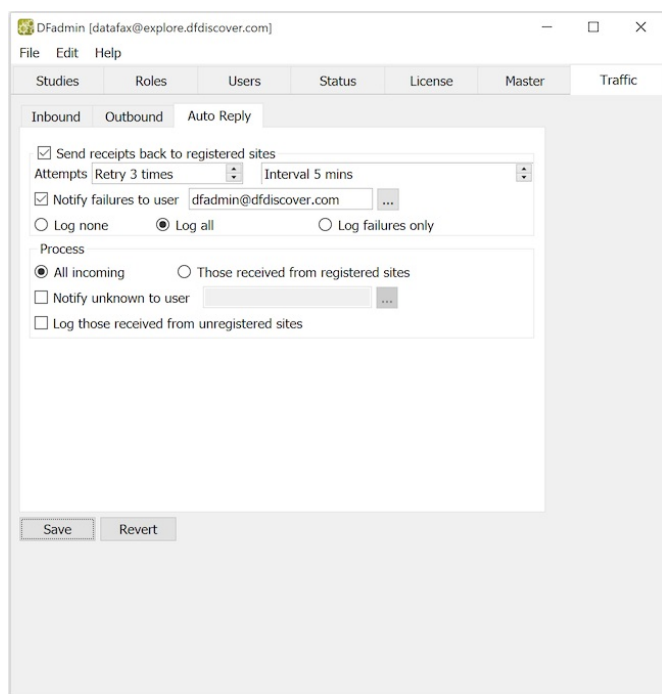
To change the status of the Outbound service, select one of:

- **Action** > **Start Outbound** Start the currently stopped Outbound service.
- **Action** > **Stop Outbound** Stop the currently processing Outbound service. Documents that are currently being transmitted will complete transmission but documents waiting in the queue will not be processed until the Outbound service is started again.

Auto Reply

Auto reply is a debugging feature useful for diagnosing problems with sites that have ongoing difficulty submitting CRFs to **DFdiscover**. The feature works by sending a receipt to selected sites when documents from them arrive into **DFdiscover**. The receipt contains information about when the document was received, how many pages it contained, etc. A site that is having difficulty sending can then wait for their Auto Reply receipt before sending more CRFs - this will tell them immediately if there is a problem.

Auto Reply is configured in **DFAdmin's Auto Reply** view.



The configuration parameters that can be altered are:

- **Send receipts back to registered sites:** Enabling this checkbox enables basic Auto Reply processing. When enabled, **DFdiscover** sends a confirmation of receipt for each document received from a registered site. Registered sites require a unique sender id for identification by **DFdiscover**. Hence, receipts are not sent to all sites that submit documents, but only to those that are registered. The remaining options are relevant only with this setting of the option.

With this option is not checked, Auto Reply is turned off and receipts are never sent to any site, whether a registered site or not.

- **Attempts:** When sending receipts, how many attempts should be made and at what interval?
- **Notify failures to user:** If an Auto Reply receipt cannot be sent, should a notification be sent to an internal administrator, and if so, to whom?
- **Log:** This is a choice field. When set to **Log none**, no log will be kept of receipts that are sent back. This log is not necessary and nothing in **DFdiscover** requires or references it; its purpose is purely informational. Setting this option to **Log all** configures **DFdiscover** to log every attempt to send a receipt. To log only those receipts that could not be sent, set this option to **Log failures only**.
- **Process:** This is a choice field that optionally allows **DFdiscover** to filter documents that originate from an unknown sender. When set to **All incoming** **DFdiscover** uses normal Inbound processing, where all received documents are processed. If the option is set to **Those received from registered sites**, only documents that arrive from registered sites are processed, all other documents from unregistered sites are discarded. This latter setting must be used with great care, as documents from unregistered sites will be deleted without further warning.
- **Notify unknown to user:** This setting applies only when **All incoming** is also set. If a document is received from an unregistered site, should a notification be sent to an internal administrator, and if so, to whom?
- **Log those received from unregistered sites:** This setting applies only when **All incoming** is also set. If a document is received from an unregistered site, should a log record be written? Log records are written to *sitefax_log*, which is described in [sitefax_log](#).

After making any changes, select **File** > **Save** to save the changes and re-configure the Auto Reply service. The changes will be applied starting with the next incoming document that is received.

Registered Sites Database

Prior to enabling auto reply processing by setting, a database of registered sites must be initiated. It does not need to be complete at this time. Since auto reply processing is based on the sender identification, the registered sites database is implemented at the generic **DFdiscover** level, not at the study specific level. At the study specific level, some of the same sites will appear in the studies database. This is expected, noting that the information required for the study sites database (which identifies where Query Reports will be sent) is different information than that required for the registered sites database (which identifies from where documents are sent).

A registered sites database must be created directly on the server. No special editor is included in **DFdiscover** for the creation or modification of this file, however, it is easily updated with a simple text editor. New sites can be added, existing sites modified, and sites deleted while the auto reply feature is enabled.

It is a plain text file, named `/opt/dfdiscover/lib/DFsites.db`, that identifies the clinical sites that will receive fax or email receipts. The format of this file is documented in [DFsites.db](#).

Typical records from *DFsites.db*

```
1|General Hospital|905-999-9876|1-905-999-9876|/opt/studies/xyz/replyfax/reply.001
1|St. Elsewhere|elsewhere@hospital.com|mailto:person@hospital.com|/opt/studies/abc/replyfax/reply.001
```

For the successful operation of auto reply and the registered sites database, each site must transmit a unique sender identification string. This can be a fax number or email address. The replyfax field (field 4) of each *DFsites.db* record may also consist of a fax number or email address. Comparison between the sender identification and the replyfax field is always performed case-insensitive.

Sites which do not need receipts are not required to be included in this file.

Auto reply receipt for registered sites

The 5th field of each site registration record identifies the name of a text file to be used as a template for that site's image receipt. The same receipt file may be shared across sites simply by referencing the same file in the 5th field of each site record that shares the file. Conversely, each site might use or require a separate receipt, for example, to allow for different languages.

The requirements for a document receipt are documented in [DFreceipt.sample](#). An example, shown below which illustrates a document receipt.

Example Auto Reply receipt file

```
DFDISCOVER RECEIPT FOR STUDY XYZ
To: %s
From: XYZ STUDY coordinating site
Re: recently received document from %g

A %p page document was received at %t on %d.

If this does not agree with your submission log, please contact:
Jane Wright at 905-999-8888, extension 7777

THANK YOU FOR YOUR CONTINUING SUPPORT
```

The symbols %s, %g, etc are placeholders for information that is filled in by **DFdiscover** during the composition of the receipt for an individual image.

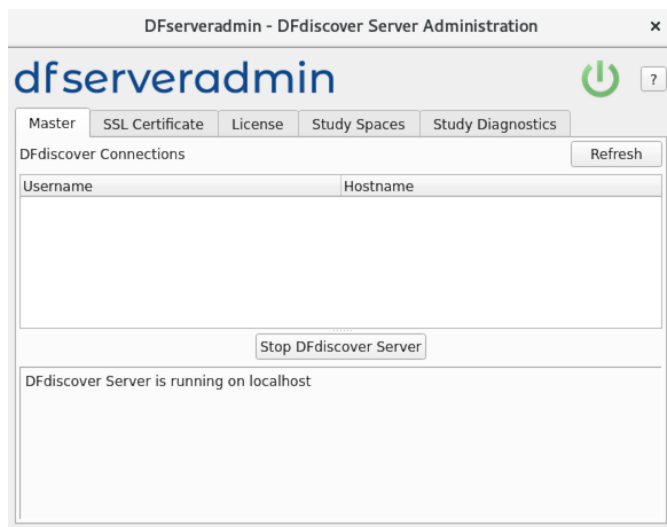
DFserveradmin

DFserveradmin is a server-based application that is used to:

- start and stop the **DFdiscover** services.
- request SSL certificate signing from DF/Net Research, Inc. (DFnet).
- manage your **DFdiscover** license and see usage statistics.
- run study diagnostics.
- check and fix study file permissions.

DFserveradmin must be started from the command-line by the system `root`. There are no options required or recognized by this application.

Invoking **DFserveradmin** at the command-line displays the main user interface for the application. If **DFdiscover** is currently running the green start icon is displayed in the header. The stop icon is displayed if **DFdiscover** is not running.



The functionality of **DFserveradmin** is grouped into several tabbed windows. Related functions appear in the same tabbed window. The purpose of each window appears in the name of the tab. Click the corresponding tab name to access its tabbed view and functionality.

Master

DFserveradmin opens with the **Master** view visible.

On an operational system, any connections to the **DFdiscover** server are listed. It is not recommended that you stop **DFdiscover** with live connections, as this will disconnect these sessions without warning the user that made the connection. The list of users is not updated in real time. Click **Refresh** to refresh the user list with current connections.

Click **Start DFdiscover Server** to start **DFdiscover**. Click **Stop DFdiscover Server** to stop **DFdiscover**. Progress messages are logged to the messages panel. Additional startup information may also be logged to `/tmp/dfstartlog`.

Request an SSL certificate signing

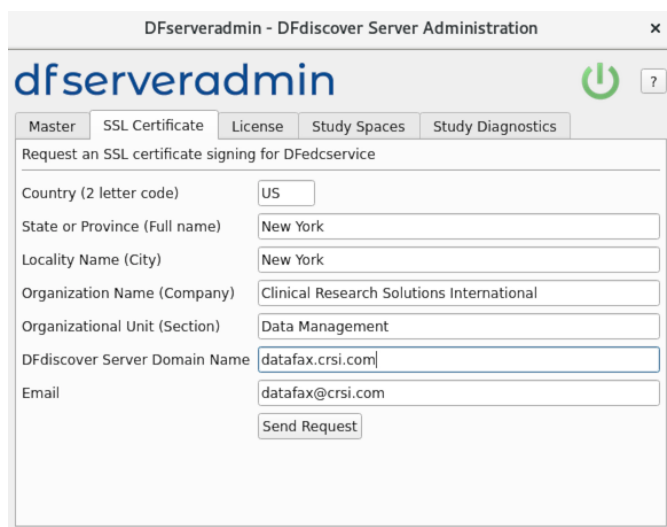
Client applications cannot connect to your **DFdiscover** server unless a signed, and current, SSL certificate is in place on the server. Signed SSL certificates are typically valid for a period of 1 or 2 years.

When installed, a signed SSL certificate tells any client applications attempting a connection that the server is known, can be trusted and that communication is encrypted.

There are several commercial certificate signing authorities. Some clients prefer to have a signed certificate from one of these authorities. For these clients, a commercial certificate can be acquired independently and installed, as described in [Installing a Signed Certificate](#).

Alternatively, DFnet can act as the signing authority for your server certificate. There is no additional charge for this service.

To use DFnet as your certificate signing authority, click the **SSL Certificate** tab.



A new system does not have a signed certificate, so this dialog must be completed and submitted for signing to DFnet. Responses are required for all fields. The responses are visible to any user that selects **Help** > **Certificate Info** in their client application. Further, the **DFdiscover Server Domain Name** value must exactly match the fully qualified domain name (FQDN) of your server. Otherwise, users will be presented with a certificate warning dialog each time that they login to your server.

Click [Send Request](#) to email the unsigned SSL certificate request to DFnet for processing.

You can expect an email response containing a signed certificate within 24 hours on business days. In the interim, it is still possible to proceed as **DFserveradmin** also generates a self-signed certificate that is valid for 10 days. The signed certificate will be emailed back to the email address that was specified as the *Email Address* in the certificate request procedure. Typically it will appear as a small (~2KB) attachment named *edccert.pem*.

NOTE: Outgoing email

The certificate signing request is emailed to DF/Net Research, Inc. from the machine where the **DFserveradmin** application was executed. This machine must be able to send email to an external recipient (DFnet). If that is not possible, the files */tmp/cert.csr.txt* and */tmp/cert.csr* can be transferred (use binary mode if doing this via ftp) to another machine and sent as attachments in an email to certreq@dfnetresearch.com.

Installing a Signed Certificate

The signed certificate, whether received from DFnet or from a commercial signing authority, must be installed in the location referenced by the *certificate* keyword in the */opt/dfdiscover/lib/DFedcservice.cf* file. For example, this *DFedcservice.cf* file:

```
master localhost
key /opt/dfdiscover/lib/edckey.pem
certificate /opt/dfdiscover/lib/edccert.pem
documentroot /opt/dfdiscover/edcdocs
```


identifies */opt/dfdiscover/lib/edccert.pem* as the location of the server certificate.

To install a new, or update an existing, certificate, save the signed certificate to this location.

A signed certificate is typically valid for one or two years and must be renewed on an annual basis.

License

Before starting the **DFdiscover** server, a valid license must be installed. Click the **License** tab to display the License view.




The screenshot shows the 'DFserveradmin - DFdiscover Server Administration' window. The 'License' tab is selected, and the 'License Settings' section is visible. The settings include:

- Host Identifier: DRHG-LPSD-LVLV-C28F-WCXT (with an Update button)
- Total Seats: 5
- Expiration Date: 2022/01/25
- Features: DFWS,LS (with a menu icon)
- Password: TGH8-VF2T-4E9Z-XE3E-9UZ6
- Hostname: localhost
- Admin Seats: 2 (with a User Seats field set to 3)

Buttons at the bottom include 'Check Validity', 'Apply', 'Revert', and 'Request New License'.

On new installations, click [Request New License](#) to display the license request form.



Request New License

License Information

Name:

Company:

Email Password to:

Email Invoice to: same as Email Password

Purchase Order #:

Total Seats (All Servers): 5 Current: 5

Total Number of Servers: Current: 1

Details for this Server

OS Version: Linux 4.18.0-294.el8.x86_64 x86_64

DFdiscover Version: 5.4.0

Number of Seats: 5

Number of Admin Seats:

License Period: 2022/01/25 to 2023/01/25

Hostname: localhost

Host Identifier: DRHG-LPSD-LVLV-C28F-WCXT

Features: DFWS.LS


Message:

Complete the form, editing any pre-filled fields as required. Click to submit the license request to DFnet for processing. A server password will be emailed back to the address provided.

NOTE: Issuing a license is not an "automatic" email response, and is subject to review and verification by DFnet staff. Normally you can expect a response with 24 hours if your request is made during technical support hours.

Study Spaces

Before new studies can be registered in **DFadmin**, the **UNIX** directories that serve as the root for **DFdiscover** study directories must be created from the **UNIX** command-line and then registered as **DFdiscover Study Spaces**.



DFserveradmin - DFdiscover Server Administration

dfserveradmin

Master | SSL Certificate | License | **Study Spaces** | Study Diagnostics

Study Spaces (directories)

Name	Directory	Status
opt_studies	/opt/studies	Active

Study Space

Name:

Directory:

Active:

When adding a new study in the **DFadmin** Studies dialog, the administrator then selects one of the predefined study spaces and enters only a study folder name. This ensures that administrators can only create study directories in locations that have been approved by the **UNIX** administrator.

NOTE: No part of a study space can itself be a **DFdiscover** study directory. For example, `/opt/studies` and `/opt/studies/ABC_Studies` could both be study spaces for the following study directories: `/opt/studies/A`, `/opt/studies/B`, `/opt/studies/NIH_Studies/A`, `/opt/studies/NIH_Studies/B`, etc., but `/opt`, `/opt/studies` and `/opt/studies/NIH_Studies` cannot be **DFdiscover** study directories.

Study spaces and directories must be unique, case-insensitive and cannot include the follow characters: `\`' $; & * < | > |` and SPACE

Study Diagnostics

If you are having problems connecting to a study, this utility may help in troubleshooting the problem. This function is equivalent to running

`/opt/dfdiscover/utis/DFstudyDiag` on your study at the command-line in a terminal session.

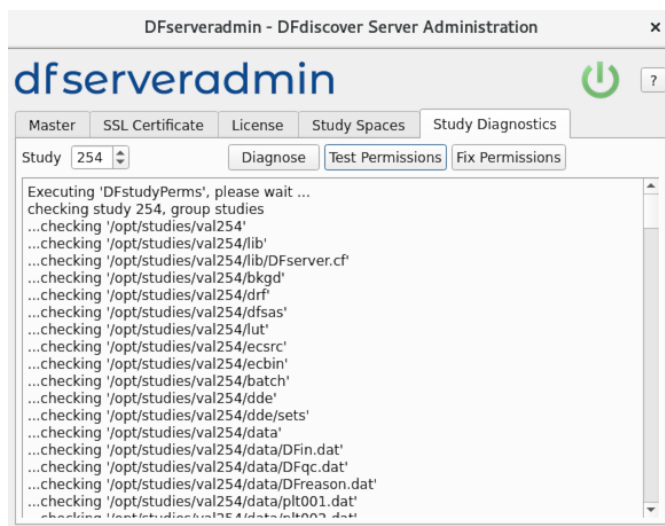
Select a study number and click **Diagnose**. The output is displayed in the output window. For more information, see [Programmer Guide, DFstudyDiag](#).



Study Permissions

If you suspect you are having problems with file-level permissions, this utility may help in troubleshooting and fixing the problem. This function is equivalent to running `/opt/dfdiscover/utis/DFstudyPerms` on your study at the command-line in a terminal session.

Select a study number and click either **Test Permissions** to troubleshoot, or **Fix Permissions** to fix any problems. Any problems or corrections are listed in the output window. For more information, see [Programmer Guide, DFstudyPerms](#).



Periodic Maintenance

This chapter discusses various topics relevant to ongoing, periodic maintenance of a **DFdiscover** system. Familiarity with these topics is recommended, and as many of the concepts as possible should be implemented in each **DFdiscover** environment. Some topics could even be paraphrased into standard operating procedures.

The most critical area of maintenance for any **DFdiscover** installation is disk usage monitoring. **DFdiscover** can be brought to a grinding halt by failing to notice that you are about to run out of disk space. Periodic monitoring of current disk usage and planning to make sure that disk space will be available for incoming images and attachments is essential to the health of the system.

Disk Maintenance

Disk usage by a **DFdiscover** study is highest for CRF image files. Generally, 90% of the disk space required for a study is consumed by CRF image files. This section covers managing disk space for the individual CRF image files that are kept in the `PAGE_DIR` directory defined for each study. [Archive File Maintenance](#) discusses management of the original TIFF/PDF files that the image files are extracted from.

Monitoring and planning future disk needs

True backups require a quiet filesystem. If this is done during off-hours, it is unlikely that any user will be modifying a study database. However CRF images may still be arriving, especially from international sites. This problem can be solved at two levels:

- disable **HylaFAX** from answering any new incoming fax calls, or
- shutdown **DFdiscover**.

The former solution has the problem that it may frustrate sites that are trying to submit CRF images while the backup is being done. The latter solution should be careful to not backup the incoming images directory as its contents will not be quiet as long as one or more incoming images are being received.

DFdiscover can be halted from a shell script (like the **cron** process that performs the backups) by executing the command.

```
# /opt/dfdiscover/bin/DFshutdown -f
```

and then subsequently restarted using

```
# /opt/dfdiscover/bin/DFbootstrap
```

If only an individual study directory (or group of study directories) needs to be backed up, then **DFdisable.rpc** is used to temporarily disable the study servers that must not be running while the backup is executing. This allows users of other studies to continue with their **DFdiscover** activities while the backup proceeds. When the backup is complete, **DFenable.rpc** must be used to re-enable the halted study servers.

There are many freely and commercially available applications for performing system backup. Discuss the options with your corporate IT team.

What should a regular **DFdiscover** backup include?

It is essential to back up the **DFdiscover** setup and configuration as well as the individual **DFdiscover** study definitions and data.

DFdiscover Setup and Configuration Files

There are two files in your `/opt/dfdiscover/lib` directory which contain information that will be essential to rebuilding your **DFdiscover** server.

- **DFstudies.db** This file contains the study number and `$STUDY_DIR` location for each study. While you may need to put your study information in a different location on a new server, this information is essential to knowing where to look for it in available backups.
- **DFstudyspaces.db** This file contains a list of the filesystem locations on a server where **DFdiscover** study information and data can be stored. As with `DFstudies.db`, knowing how the old system was configured is important even if the new system needs to be different for some reason.

There are also non-study-specific files that **DFdiscover** updates regularly, and hence they also need to be part of a regular backup. These include the directories:

- `/opt/dfdiscover/work`
- `/opt/dfdiscover/lut`
- `/opt/dfdiscover/ecsrc`
- `/opt/dfdiscover/ecbin`
- `/opt/dfdiscover/lib`
- `/opt/dfdiscover/archive`, or whatever the local setting is for the TIFF/PDF archive directory

Study Files

This includes the **DFdiscover** setup and configuration files, the data records themselves, and the images, if any, associated with them. At a minimum, the directories to backup include:

- `$STUDY_DIR/bkgd`
- `$STUDY_DIR/data`
- `$STUDY_DIR/dfsas`
- `$STUDY_DIR/dfschema`
- `$STUDY_DIR/drf`
- `$STUDY_DIR/ecbin`
- `$STUDY_DIR/ecsrc`
- `$STUDY_DIR/lib`
- `$STUDY_DIR/lut`
- `$STUDY_DIR/pages`
- `$STUDY_DIR/pages_hd`
- `$STUDY_DIR/reports/QC`
- `$STUDY_DIR/work`

NOTE: The `pages_hd` directory may not be present in studies where HD resolution has not been enabled.

There may also be other directories used by a study team that are not specific to, or required by, **DFdiscover** - consult with the **DFdiscover** users to identify what those directories or files might be. It may be safest, and most inclusive, to specify the study root directory for backup. In this way every sub-directory will by default be included.

Use of tar to backup a study

```
# cd /opt/studies/val254
# tar cf /dev/rmt/0 bkgd data dfsas dfschema drf ecbin ecsrc lib lut pages pages_hd reports/QC work
```

In this case, the study is rooted at `/opt/studies/val254`

If there are other directories or files to be included, it may be safest to capture the entire study hierarchy with this command:

```
# cd /opt/studies/val254
# tar cf /dev/rmt/0 .
```

or this command:

```
# cd /opt/studies/
# tar cf /dev/rmt/0 val254
```

The former excludes the study parent directory name from the backup, while the latter includes it.

Testing

As with any backup or disaster recovery solution, it must be tested to confirm that it is operating in the expected manner, that all of the needed contents are in fact being backed up and that is occurring on the planned, regular schedule.

DF/Net Research, Inc. encourages all clients to test their backups at least once per month. Additional, manual backups should be performed for "milestone" events - for example, launch of a new study, close of a completed study, or before upgrading to a new software version. Individual needs and resources will vary. Clients are also encouraged to have a secondary, standby server available at all times and to regularly update the secondary server with the contents of the primary server.

Archive File Maintenance

This section covers the management of the original archive (TIFF or PDF) files that the PNG files are extracted from. In **DFdiscover**, these files contain the original image transmission as received from the sender via the fax modem or scan transmission via **DFsend**. Each file contains the total number of pages sent in the transmission.

Archiving of TIFF and PDF files is controlled by the value of the `INBOUND_ARCHIVE_DURATION` variable defined in the configuration of incoming daemon(s). Unless archiving has been explicitly disabled by setting the value of this parameter to 0, each incoming file is archived by the **DFdiscover** incoming daemon. The value of the archived image is in the ability of an administrator to subsequently manually recover pages that are mistakenly deleted. This process is described in [Retrieving lost CRF images](#). The **DFdiscover** software itself does not require archived images, nor does it confirm their existence.

Generally speaking, archive files should be routinely saved to secondary media (tape) and then deleted from primary storage (disk). How many archive files are kept on disk before being moved to secondary storage is a matter of individual preference and comfort level, but an average of 4 weeks of archived images is appropriate. This leads to a monthly procedure in which any archive files that are more than 4 weeks old are moved to secondary storage and deleted from disk.

By way of example, consider an environment where the archive files have never been moved to secondary storage and it is now desired to begin implementing a routine monthly procedure for doing this. The archive files have been kept in `/opt/dfdiscover/archive` and all but the four most recent weeks worth of files must be moved to tape storage on device `/dev/rmt/0`. The archive directory has the following contents:

```
# ls /opt/dfdiscover/archive
1716/  1720/  1724/  1728/
1717/  1721/  1725/  1729/
1718/  1722/  1726/  1730/
1719/  1723/  1727/  1731/
```

The following command would have the desired result of archiving the oldest weeks to tape:

```
# cd /opt/dfdiscover/archive
# tar cvf /dev/rmt/0 171{6,7,8,9} 132{0,1,2,3,4,5,6,7} | lp
```

The **tar** command is used to backup the files to tape in this case, but other backup commands are equally valid. This particular **tar** command also creates a table of contents listing as the backup is created, and that listing is directed to the default printer. This provides a convenient, printed table of contents that can be kept with the tape.

The next step is to delete from primary storage the archive files that have been copied to secondary storage. Before deleting the archive files, verify that the backup created on secondary media is complete. This confirmation can be done with a visual review of the printed table of contents or by immediately performing a test restore of the media to another location.

After confirming that the backup copy of the TIFF files is complete, delete the original copies of the files from disk. Continuing with the example, the command to execute is:

```
# cd /opt/dfdiscover/archive
# /bin/rm -rf 171{6,7,8,9} 132{0,1,2,3,4,5,6,7}
```

Having completed the steps, the directory would have the following contents:

```
# ls /opt/dfdiscover/archive
1728/  1729/  1730/  1731/
```

The only remaining step is to formalize this process into a periodic routine.

Study Maintenance

This section describes those activities that should be executed on a regular database as part of a pro-active study maintenance process.

Before going 'live' with a new study setup, it is advisable to restore the study directory to a base (empty) state that does not contain any test data or test images.

Periodically the study directories should be examined and old, stale files removed. These files are typically temporary files that were created by users in the study work directory and quality control reports that were created but never sent.

It may also be required to perform regular archiving of a study database for interim analyses.

Going live with a DFdiscover study

It is highly recommended that a study setup be thoroughly tested before real data is accepted from investigative sites. For paper-based studies, this testing should include completing blank case report forms with actual data, submitting the case report forms into the system, validating the ICRed data records, and creating Query Reports to test the visit map, page map, and/or conditional plate and termination maps. The **DFdiscover Study Setup Worksheets** are an excellent aid in ensuring that all of the required steps are completed and documented.

The result of this testing will be a study database that contains data, CRF images, and Query Reports that are not relevant to the real study data. It is important before going live with a study to remove all of this test data. It is straightforward to remove this test data before the real data arrives; it is much more tedious to remove it once it becomes combined with real subject data.

Deleting test data

To delete the existing test data from a study the following steps should be followed.

1. Disable the study server

It is a requirement that the study server be disabled when the test data is deleted. This can be done either via the **Status** dialog of **DFadmin** or from the command-line, using **DFdisable.rpc**, as illustrated in following example.

Disabling study 254:

```
# /opt/dfdiscover/bin/DFdisable.rpc -s 254 "going live"
```

2. Remove, or rename, the existing *data pages pages_hd* and *reports/QC* directories

Note that removal of these directories assumes that they contained only information that was created by **DFdiscover**. If these directories contain other information that is outside the control of **DFdiscover** (and this is not recommended), then they cannot simply be deleted.

Removing the directories containing test data for study 254:

```
# cd /opt/studies/val254
# rm -r data pages pages_hd reports/QC
```

3. Enable the study server

Once the study server starts again, the removed directories will be re-created as empty directories.

Enabling study 254:

```
# /opt/dfdiscover/bin/DFenable.rpc -s 254 "going live"
```

For documentation purposes, the setup should be printed from the **[Study] > [Print]** menu in **DFsetup** and the current user permissions should be printed from the **Permissions** dialog in **DFadmin**.

Monitoring study directory permissions

A **DFdiscover** study is stored on disk as an inverted tree structure in the filesystem. The information required at any moment during the use of a study is available as one or more files in that filesystem structure. Using the **UNIX** filesystem directly has the advantage that this same information is also readily available to applications outside of **DFdiscover**, for example, for the purposes of scripting or working with third-party applications. However, this flexibility also has the drawback that **UNIX** filesystem permissions and the permissions required by **DFdiscover** are not always in perfect agreement. This can lead to users that are unable to open files that should otherwise be permitted to. The purpose of this section is to describe the permissions that **DFdiscover** requires and suggest ongoing maintenance to ensure that those permissions are maintained.

Owner and group

By default, **DFdiscover** will create all of the needed directories and files for a study with owner *datafax* and group *studies*. The ownership should always remain as *datafax*. The group *studies* is intended for general sharing of study files across all **DFdiscover** users. This typically matches the primary group assigned during login to **DFdiscover** user accounts. If a different group is being used for the study, then that group name should be applied to all of the directories and files. At the same time, that group name should be listed as the primary group for login to those **DFdiscover** accounts that are specific to the study.

No permissions are required for *other*, and so they are not granted by **DFdiscover**. It should be possible to accomplish all needed tasks with owner or group permissions.

Owner and group settings are not applied by **DFdiscover** to directories or files which it does not create. For example, a *sas* or *batch* sub-directory, which is created by a user will not have the same ownership and group. It is recommended that owner *datafax* and group *studies* be applied to these directories and files, but this must be done manually.

Maintaining study filesystem permissions

DFdiscover includes a utility application, **DFstudyPerms**, (see [Programmer Guide, DFstudyPerms](#)) which examines, reports, and optionally repairs permissions for a study filesystem. This application should be run from the command-line whenever a permissions problem is suspected and also as part of a regular maintenance procedure to identify and correct problems with permissions.

To report on study permission problems, any user can execute the command:

```
% /opt/dfdiscover/utills/DFstudyPerms #
```

where # is the study number. Run in this fashion, **DFstudyPerms** remains silent unless a problem is discovered. Any permissions which do not match the expected permissions are reported, one line per file or directory. It also uses the group *studies* unless another group is specified with the *-g groupname* option.

To fix study permission problems, the *root* account is required. In this case the command is:

```
# /opt/dfdiscover/utills/DFstudyPerms -f #
```

where # is the study number and *-f* instructs the application to correct any permission errors that it encounters. Again, the *-g groupname* option is needed if the study group is not *studies*.

It is recommended that the latter invocation be added to *root's crontab* and executed at least once per month.

Default study filesystem permissions

Study filesystem permissions table below lists the study filesystem permissions. The permissions are reported as 3 triples of 3 characters. The first triple is owner permissions, the second group, and the third other. The 3 character positions, *rx*, represent read permission, write permission, and search permission respectively. If a particular permission is not granted, it appears as a dash, *-*, in the listing. If a file or directory is checked by **DFstudyPerms** it is also checked for either exact permissions or minimum permissions. If it is checked for exact permissions, it must have exactly the listed permissions - any other permission will generate a message. If it is checked for minimum permissions, then additional permissions (for example, additional write permissions for group) are acceptable and will not generate a message.

NOTE: Most of the permissions are checked by **DFstudyPerms** but not all of them. It is expected that a future version of **DFstudyPerms** will include checking of these additional files.

Study filesystem permissions

Name	File or Directory	Permissions	Type of check	Notes
.	Directory	<i>rxr-x-</i>	Minimum	This is the study parent directory. If users are permitted to create their own sub-directories, the permissions will need to be <i>rxrwx-</i>
<i>batch</i>	Directory	<i>rxr-x-</i>	Minimum	
<i>bkgd</i>	Directory	<i>rxrwx-</i>	Minimum	
<i>bkgd/DFbkgd???.tif</i>	File	<i>rw-rw-</i>	Minimum	
<i>bkgd/plt???</i>	File	<i>rw-rw-</i>	Minimum	
<i>bkgd/DFbkgd???</i>	File	<i>rw-rw-</i>	Minimum	
<i>data</i>	Directory	<i>rxr-x-</i>	Minimum	Write permissions on this directory should never be granted to any account other than <i>datafax</i> .
<i>data/.dat</i>	File	<i>rw---</i>	Exact	
<i>data/.idx</i>	File	<i>rw---</i>	Exact	
<i>data/.jnl</i>	File	<i>rw-r-</i>	Exact	These audit trail files must not be writable by any account other than <i>datafax</i> . They are readable for the purposes of audit trail reports like DF_ATmods .
<i>drf</i>	Directory	<i>rxrwx-</i>	Minimum	
<i>dde</i>	Directory	<i>rxrwx-</i>	Minimum	
<i>dde/sets</i>	Directory	<i>rxrwx-</i>	Minimum	
<i>dfsas</i>	Directory	<i>rxrwx-</i>	Minimum	
<i>ecbin</i>	Directory	<i>rxr-x-</i>	Minimum	
<i>ecsrc</i>	Directory	<i>rxr-x-</i>	Minimum	

Name	File or Directory	Permissions	Type of check	Notes
<i>lib</i>	Directory	<i>rw-rwx---</i>	Minimum	
<i>lib/DFcenters</i>	File	<i>rw-rw---</i>	Minimum	
<i>lib/DFfile_map</i>	File	<i>rw-rw---</i>	Minimum	
<i>lib/DFschema</i>	File	<i>rw-rw---</i>	Minimum	
<i>lib/DFschema.stl</i>	File	<i>rw-rw---</i>	Minimum	
<i>lib/DFserver.cf</i>	File	<i>rw-r---</i>	Exact	
<i>lib/DFsetup</i>	File	<i>rw-rw---</i>	Minimum	
<i>lib/DFsetup.backup</i>	File	<i>rw-rw---</i>	Minimum	This file contains the previous version of the study setup and is overwritten as part of the initialization process of DFsetup .
<i>lib/DFtips</i>	File	<i>rw-rw---</i>	Minimum	
<i>lib/DFvisit_map</i>	File	<i>rw-rw---</i>	Minimum	
<i>lib/DFccycle_map</i>	File	<i>rw-rw---</i>	Minimum	These remaining files in the study <i>lib</i> directory are optional.
<i>lib/DFcplate_map</i>	File	<i>rw-rw---</i>	Minimum	
<i>lib/DFcterm_map</i>	File	<i>rw-rw---</i>	Minimum	
<i>lib/DFcvisit_map</i>	File	<i>rw-rw---</i>	Minimum	
<i>lib/DFedits</i>	File	<i>rw-rw---</i>	Minimum	
<i>lib/DFlut_map</i>	File	<i>rw-rw---</i>	Minimum	
<i>lib/DFmissing_map</i>	File	<i>rw-rw---</i>	Minimum	
<i>lib/DFpage_map</i>	File	<i>rw-rw---</i>	Minimum	
<i>lib/DFqcproblem_map</i>	File	<i>rw-r---</i>	Minimum	
<i>lib/DFqcps.prolog</i>	File	<i>r-r---</i>	Minimum	
<i>lib/DFqcsort</i>	File	<i>rw-rw---</i>	Minimum	
<i>lib/DFraw_map</i>	File	<i>rw-rw---</i>	Minimum	
<i>lib/QCcovers</i>	File	<i>rw-rw---</i>	Minimum	
<i>lib/QCmessages</i>	File	<i>rw-rw---</i>	Minimum	
<i>lib/QCtitles</i>	File	<i>rw-rw---</i>	Minimum	
<i>lut</i>	Directory	<i>rw-r-x---</i>	Minimum	
<i>pages, pages_hd</i>	Directory	<i>rw-r-x---</i>	Minimum	
<i>pages/????, pages_hd/????</i>	Directory	<i>rw-r-x---</i>	Minimum	These are the directories, organized by year and week of year, in which the CRF images are stored.
<i>pages/????/???????, pages_hd/????/???????</i>	File	<i>rw-r---</i>	Exact	
<i>reports</i>	Directory	<i>rw-r-x---</i>	Minimum	If users are permitted to install their own study-specific reports, these permissions will need to be <i>rw-rwx---</i> .
<i>reports/QC</i>	Directory	<i>rw-rws---</i>	Minimum	
<i>reports/QC/-???????</i>	File	<i>rw-rw---</i>	Minimum	
<i>reports/QC/QC_LOG</i>	File	<i>rw-rw---</i>	Minimum	
<i>reports/QC/QC_NEW</i>	File	<i>rw-rw---</i>	Minimum	

Name	File or Directory	Permissions	Type of check	Notes
<i>reports/QC/SENFAX.log</i>	File	<i>rw-rw---</i>	Minimum	
<i>reports/QC/SENFAX.qup</i>	File	<i>rw-rw---</i>	Minimum	
<i>reports/QC/internal</i>	Directory	<i>rxwxrwx---</i>	Minimum	
<i>reports/QC/sent</i>	Directory	<i>rxwxrwx---</i>	Minimum	
<i>reports/QC/sent/-??????</i>	File	<i>rw-rw---</i>	Minimum	
<i>work</i>	Directory	<i>rxwxrwx---</i>	Minimum	
<i>work/DFvisit.dates</i>	File	<i>rw-rw---</i>	Minimum	
<i>work/DFX_*</i>	File	<i>rw-rw---</i>	Minimum	
<i>work/DF.drf</i>	File	<i>rw-rw---</i>	Minimum	
<i>work/DF_QCupdate.log</i>	File	<i>rw-rw---</i>	Minimum	

Cleaning the study *work* directory

The *work* directory for a **DFdiscover** study includes a mixture of temporary files created by **DFdiscover** and temporary files created by users.

Files that have names beginning with *DFX* are created by the **DFdiscover DF_XXkeys** report. They are overwritten each time that **DF_XXkeys** or **DF_QCupdate** is executed. In most circumstances, they should be left alone. However, if disk space is at a premium they can be deleted, as they will be re-created the next time the reports are run.

The other temporary files that might be found in the work directory will be specific to each **DFdiscover** installation. You will have to use your own discretion in deciding which files to delete. As a general guideline, files with the following attributes are good candidates for deletion:

- at least one month old,
- created by a user other than user *datafax*, and
- have typical temporary file names like *temp*, *tmp*, *test*, and *NoName*

Cleaning the study Query Reports directory

The reports directory for a **DFdiscover** study includes study specific reports as well as Query Reports. The Query Reports are stored in a further *QC* sub-directory of the *reports* directory. Reports that are created by **DF_QCreports** are stored in this *QC* sub-directory and then are moved to a further *QC/sent* sub-directory when they are successfully sent to investigators. If reports are created by **DF_QCreports** but are never subsequently sent, they will be left in the *QC* sub-directory. Periodically check the files in the *QC* sub-directory of the study *reports* directories for such reports and delete them if they are out of date. If there is any doubt, this step should be coordinated with the staff member responsible for creating Query Reports for the study.

Procedures for archiving a **DFdiscover** study

When preparing to close-out a study or archive a copy for interim analysis, the following issues need to be considered:

- **The current state of the study setup needs to be archived.** All of this setup information is, under normal circumstances, in the study *lib* directory. However, lookup tables, for example, may reside elsewhere.
- **Is the new record queue empty?** Ideally, there should be no new records awaiting validation.
- **What data needs to be archived?** Does all of the data need to be archived? Primary records only? Are the journal files also required?
- **Do the CRF images need to be archived?** Almost always, the answer to this question is yes. The CRF images must be archived but is unlikely that there will be sufficient primary (disk) storage available to maintain an archive copy. Hence the CRF images should be archived to tape, DVD or cloud storage. The requirements for keeping the CRF images can be quite onerous and hence it is important to choose a secondary storage medium that will be readable many years in the future.
- **If the study is being closed out, DFdiscover permissions should be revoked for all users that have access to the study.** Minimally, each previously permitted user should be assigned a role that permits view-only, and eventually permissions should be completely removed.
- **Disable or de-register the study.** The study may also be disabled, so that no users can access it, or deleted from the **DFdiscover** studies database. The latter solution is ultimately preferred as this guarantees that **DFdiscover** will not process incoming images for the study to the study new queue.

A minimal set of steps for making an archive copy of the primary database records might follow the scenario outlined below.

Making an archive copy of the primary database records for study 254:

```
# mkdir -p /opt/archive254/data
# cd /opt/studies/val254
# tar cf - lib | ( cd /opt/archive254; tar xpf - )
# foreach p ( `DFlistplates.rpc -s 254` 0 511 )
? DFexport.rpc -s primary 254 $p /opt/archive254/data/exp$p
? end
# tar cf /dev/rmt/0 pages
```

This example makes an archive of the study setup and primary data for study 254 in a separate archive directory, /opt/archive254. Additionally, all of the current CRF images are archived to tape.

Retrieving lost CRF images

Rarely, a user may encounter the message 'image not available' in the **DFexplore Review Images** dialog. Before retrieving a lost image file, attempt to determine the cause of the problem and log it. Periodically review the log to look for any systematic problems that might be correctable.

Retrieving lost CRF images

The steps to retrieving a lost CRF image are as follows.

1. Determine the name of the lost CRF image

The name of the CRF image in question will be the name that followed the **Can't load** warning in the message window. The name will begin with the study pages directory and end with digits in the form **YYWW/FFFFPPP**, where **YYWW** is the year and week that the document was received, **FFFF** is the parent document's sequential number within **YYWW**, and **PPP** is the page number within the document.

2. Determine if the lost CRF image is still in the filesystem

DFdiscover does not ever delete image files - instead the image is renamed by prepending the name with an *X* and the file permissions are set so that the file is not accessible by a typical user. Therefore, even a deleted file is still present in the filesystem. **DFdiscover** always attempts to retrieve and restore lost CRF images on its own. If **DFdiscover** is not able to do so, the following procedures should be performed by the **DFdiscover** administrator. Generally this includes undoing the file renaming and setting the permissions so that the file is again accessible. This may be needed for only the *pages* directory, or it may also be required for the *pages_hd* directory if HD imaging is enabling.

Restoring a CRF image by renaming:

Suppose that **DFexplore** reports that image '1601/0023002' is not available. Looking in the filesystem under the directory where the image should be stored, the administrator sees:

```
# cd /opt/studies/mystudy/pages/1601
# ls -l *0023002
-r----- 1 datafax studies 38951 Jan 5 10:06 X0023002
```

which confirms that the CRF image file is still present in the filesystem. To restore the CRF image then requires the steps:

```
# mv X0023002 0023002
# chmod 640 0023002
```

whereby the file name is restored by removing the leading *X* and restoring the permissions so that the file can be seen by members of the study group. This should then be repeated with the *pages_hd* directory. In this directory, it may be that:

- the same renaming is required, or
- the image is not present at all.

The latter case is not unusual - it would indicate that HD imaging was not enabled at the time that the image was first received. In such a case there would be no need to restore the HD image in the *pages_hd* directory.

DFdiscover Study Consistency

The **DFdiscover** system includes several reports that target potential problem areas in a study setup and study database. These reports are **DF_ICrecords**, **DF_ICimages**, **DF_ICqcs**, **DF_ICkeys**, **DF_ICvisitmap**, and **DF_ICvisitdates**. This section concentrates on the **DF_ICrecords**, **DF_ICimages**, and **DF_ICqcs** reports. Any failure output from these reports represents a consistency error requiring **DFdiscover** administration privileges to resolve. The remaining reports [Standard Reports Guide, DF_ICkeys](#), [Standard Reports Guide, DF_ICvisitmap](#) and [Standard Reports Guide, DF_ICvisitdates](#) detect consistency errors that a user can resolve.

DF_ICrecords

The **DF_ICrecords** report verifies the integrity of data records for all or specified plates in the database. It does this by confirming that each record has the correct number of fields defined by the plate definition in the study setup. Additionally, **DF_ICrecords** performs the following checks on each record in the specified data files:

- the record has the correct study and plate number,
- the record has properly formatted creation and modification timestamps
- there is exactly one primary record for the record's key fields

The latter check detects more than one primary record for a set of keys and also detects secondary records that have no primary.

Executing this report with the *-d* option creates a DRF named *ICrecords.drf* that contains a record for each data record that fails one or more of the above checks. Using **Select-By Data Retrieval File DFexplore** is used to correct each problem record detected by **DF_ICrecords**. After resolving the problems, re-execution of **DF_ICrecords** will generate no error output.

In addition to the **DF_ICrecords** report, the shell-level utility, **DFcmpSchema**, are used to more stringently examine each record. **DF_ICrecords** ensures that the database structure is consistent with **DFdiscover** requirements. **DFcmpSchema** ensures that the database content is consistent with the study schema.

DF_ICimages

The **DF_ICimages** report verifies that each data record in a study database references a CRF image in the study pages directory, and conversely that each CRF image in the study pages directory is referenced by exactly one data record.

In most cases, the **DF_ICimages** report should be run with the `-x` option which forces the report to execute with the database in a read-only state. Without this option, the database is in a read-write state that allows the database state to change while the report is being run. The end result may be that **DF_ICimages** indicates problems with are present because they are timing related.

If the **DF_ICimages** report detects a record that references a missing CRF image, follow the steps in [Retrieving lost CRF images](#).

If the **DF_ICimages** report detects a CRF image that is not referenced by a data record, two resolution methods are possible:

- Move the CRF image from the study pages directory to the `/opt/dfdiscover/identify` directory so that it can be re-entered into the study new queue.

For example, if **DF_ICimages** indicates that the CRF image `9901/0045001` does not have a corresponding data record, the following command will move the CRF image back to the identify directory for subsequent identification and re-processing:

```
# cd /studies/mystudy/pages
# mv 9901/0045001 /opt/dfdiscover/identify/9901.0045001
```

- Locate the original journal entry for the record in the study journal files and re-submit that (edited) journal record with **DFimport.rpc**.

Using the same example image name, the steps are to locate the original journal entry for the record (the original entry is denoted with leading text of `d/0/0`), edit the journal record, and pass the result to **DFimport.rpc**. **DFimport.rpc** requires the study number.

Restoring a record from the journal for study 254

```
# cd /studies/mystudy/data
# grep "d|0|0|9901/0045001" *.jnl | \
/opt/dfdiscover/bin/DFget 5-NF | /opt/dfdiscover/bin/DFimport.rpc -an 254 -
```

The needed steps can be accomplished with one command that locates the needed journal record (using **grep**), removes the leading 4 fields of the journal record (using **DFget**), and finally imports the record by adding it to the new record queue using **DFimport.rpc**.

Finally, if **DF_ICimages** detects a CRF image that is referenced by two or more data records, **DFexplore** is used to review all of the involved records and delete all but the correct primary (or secondary) record.

DF_ICqcs

The **DF_ICqcs** report:

- detects final database records that have one or more unresolved queries
- detects queries that are not referenced by the key fields in any data record (free floating queries)
- detects multiple queries that reference the same data field (duplicate queries)

The **DF_ICqcs** report includes the `-r` option that causes the report to attempt to repair problems resulting from un-referenced queries and final records having unresolved queries. Inconsistencies are resolved by deleting all un-referenced queries. On final records, the unresolved queries are marked as resolved.

Multiple queries that reference the same data field can be resolved by using **DFexplore** or **DFweb** to delete all but one of the duplicate queries.

System Maintenance


A **DFdiscover** system as a whole also needs routine maintenance. This maintenance includes regular, generally daily, backups of important filesystems as already described, as well as routine pruning of the filesystem that involves truncating log files.

Certificate Info

Each of the client applications communicates with the **DFdiscover** server using HTTPS on port 443. This port must be open on any firewalls between the local computer and the study server.

This is industry-standard technology that encrypts the bi-directional communication using a 'certificate of trust' provided by the server. It is the same technology used by banks and the majority of secure, global web services.

You can visually confirm that the communication is secure. After logging-in to **DFadmin** select **Help** > **Certificate Info** and look for the green checkmark.



Renewing the Certificate

In the **Certificate Info** dialog, take note of the expiry date. The certificate for your server is valid for a defined period of time.

If the certificate expires, clients will not be able to connect using encrypted communication. It is your responsibility to ensure that certificate expiry does not happen. This is easy to handle.

The certificate issuer for your **DFdiscover** server is identified in the value of the **Issued by** field. If

- DF/Net Research, Inc. is your certificate issuer, use the command-line **DFcertReq** utility or **DFserveradmin** to request a new certificate
- DF/Net Research, Inc. is not your certificate issuer, contact the certificate issuer directly to arrange a new certificate

Truncating **DFdiscover** log files

There are various log files that are maintained by **DFdiscover** that can be periodically truncated. In truncating these files it is important to maintain the file permission and ownerships that were in place before the file was truncated. Also, you should choose between completely clearing all of the log messages or maintaining a context of the most recently written log messages. In the examples below, both methods are indicated.

`/opt/dfdiscover/work/server_log`

The **DFmaster.rpcd** application appends an entry to this log file each time a study database server starts or stops. These entries are useful in debugging but are not required for the proper functioning of a **DFdiscover** system.

Entries are expected to appear in pairs and have the following appearance:

```
DFserver.rpc.251[27239]: start on teamserver at Mon Jan 22 17:23:37 2018
DFserver.rpc.251[27239]: exit at Tue Jan 23 09:29:53 2018
```

Messages may be appended to this file between the start and exit messages, but each start should eventually be terminated by an exit.

Messages are also appended to this file if a **DFdiscover** administrator deletes a study, study data, or study data and setup info using the **DFadmin** 'Delete' option. The example below illustrates the message from each of these operations performed on study 101.

```
DFedcservice.101[21877]: jack@localhost deleted all study data Fri Dec 1 12:00:09 2017
DFedcservice.101[21877]: jack@localhost deleted all study data and setup info Fri Dec 1 12:01:19 2017
DFedcservice.101[21877]: jack@localhost deleted study from datafax registry Fri Dec 1 12:02:58 2017
```

This file can be pruned at any time; the **DFmaster.rpcd** process will re-create or re-synchronize with the file after any changes. Pruning can be accomplished from the command-line as described below.

- To clear all messages:

```
# cat /dev/null > /opt/dfdiscover/work/server_log
```

- To maintain the 50 most recent messages:

```
# cd /opt/dfdiscover/work
# tail -50 server_log > new_server_log
# mv new_server_log server_log
# chown datafax:studies server_log
```

`/opt/dfdiscover/work/sendfax_log`

Each transmitted image, independent of originating study, adds a record to this file. The record includes information about the user name of the sender, the name of the transmitted file, the date and time of transmission, and the disposition status (sent/failed) of the transmission. This information is not used by any **DFdiscover** application or report and is intended to be a debugging aid in the case of failed transmissions.

This file can be pruned in the same manner as the `work/server_log` file and can be pruned at any time. Since the file does not grow very large or very quickly, it is safe to prune this file on a quarterly, semi-annual, or even annual basis.

Files to not prune

Certain log files contain information that is relevant to a **DFdiscover** installation over its entire history. These log files should *not be pruned*.

`/opt/dfdiscover/work/fax_log`

This file contains a record for each incoming image that has been received by the **DFdiscover** system, independent of destination study. Each record includes information on the name of the received image (the `YYWW/SSSS` part is particularly important), the number of pages, the sender identification, and the date and time of receipt. This information is subsequently used by the `[Page] > [Context]` option in **DFexplore**, as well as reports: **DF_ATfaxes**, **DF_WFcrfsperwk**, and **DF_XXtime**.

The contents of this file are also mirrored by an index file, `work/fax_log.idx`. The contents of these two files *must absolutely* remain in sync.

`/opt/dfdiscover/work/.seqYYWW`

The unique sequence number that belongs to an image is determined at the time of document arrival by the **DFmaster.rpcd** process. The **DFmaster.rpcd** process determines the sequence number by consulting the appropriate `.seqYYWW` file. Under normal circumstances, only the `.seqYYWW` file for the current week is required. However, should a document need to be re-processed from the TIFF/PDF archive, the `.seqYYWW` file for the original year and week of receipt will be consulted, not the `.seqYYWW` for the current year and week. As a result, it is important that these files not be removed; this is partially the reason why

they are named with a leading dot (.).

Truncating HylaFAX log files

HylaFAX provides a detailed log of all transactions that is very useful in debugging faxing problems. The information contained in these log files includes the remote fax machine number, the speed and encoding method used to transfer the fax and information about the duration and success or failure of each transmission. These log files need to be cleaned up periodically, and **HylaFAX** provides two scripts to accomplish this.

The first of these, **faxcron**, truncates the log files, and the second, **faxqclean**, is responsible for purging job description and old document files that are left over after a fax request has completed. Both of these scripts are normally run automatically by the UNIX **cron** facility.

To ensure that scripts have been correctly configured on your machine you will need to log in as root (or have a super-user perform these steps) and execute the following commands:

```
# crontab -l > mycronjobs
# more mycronjobs
```

If you see lines containing **faxqclean** and **faxcron**, the scripts are already correctly installed and no further action is necessary. If they do not appear, edit the mycronjobs file and add the following lines to the end of the file:

```
25 23 * * * /opt/hylafax/sbin/faxqclean
0 3 * * 0 /opt/hylafax/sbin/faxcron
```

which executes the **faxqclean** script every day at 11:25PM and the **faxcron** script every Sunday at 03:00AM. Save the file and then inform **cron** of the changes with the command:

```
# crontab mycronjobs
```

It may be necessary to increase the frequency of the script execution for a very high-volume site. In such a case, it can occur that the partition containing the **HylaFAX** logs (typically */var*) will fill with log files leaving no space for normal system operation.

Conclusion

Periodic maintenance of a **DFdiscover** system as described in this chapter is a preventive measure that can save many hours or days of corrective or restorative work. It also gives **DFdiscover** users a feeling of confidence that the system is always available and running smoothly. Done regularly, this maintenance should require no more than 30 to 45 minutes per week.

Troubleshooting

This chapter provides troubleshooting assistance for the **DFdiscover** administrator.

DFdiscover is a complex system. There are many applications that communicate with each other, and each application relies on not only its own health but also on the health of those applications that it communicates with. To alleviate some of this complexity when errors occur, **DFdiscover** employs the syslog facility for errors and informational messages generated by the daemon and server applications. In addition, the client applications present most of their error messages in dialog boxes as soon as the errors occur.

System Error Log

Errors, warnings and other messages, generated by **DFdiscover** daemon and server applications, are sent to syslog, which in turn writes the messages to the system log files, as configured in */etc/syslog.conf*. By default, all messages from **DFdiscover** servers and daemons with a severity level of err and higher are logged to */var/log/messages*.

In many cases the messages will not appear elsewhere, and client applications may not even be aware that an error has occurred on the server end.

As **DFdiscover** administrator, one should regularly review the system error log, particularly the end of the file, as that is where the most recent messages are.

Message Format

Most **DFdiscover** messages are written to the syslog "user" facility, with a severity level of "err" or "crit" for more serious errors. **DFdiscover** daemon messages are written to the syslog "daemon" facility. The default severity level is "notice". **DFedcservice** also issues login and study access messages at the "warning" level and extensive logging messages at the "info" level that are ignored by syslog using the default configuration.

The logged message includes a syslog message header and a message body. The message header consists of a facility indicator, a severity level indicator, a timestamp, a tag string, and optionally the process ID. The message body contains information specific to **DFdiscover**. This includes the hostname of the machine that it occurred on, and the application name that generated the error message.

The format of this header information is as follows:

```
SYSLOG_HEADER machine program[username:studynumber:PID]: message
(1) (2) (3) (4) (5) (6) (7)
```

(1)	The syslog message header, which may vary from system to system.
(2)	The hostname of the machine that the error occurred on
(3)	The name of the application that generated the message
(4)	The owner of the application instance
(5)	The study number associated with the application instance if known
(6)	The unique process identifier of the application instance
(7)	The actual text of the message

When unexpected events or errors occur in a **DFdiscover** system, the first place to start looking is in the system error log.

Login Error Messages

This section describes some common messages that users may encounter during login while using **DFdiscover** applications such as the data collection tools, **DFsetup**, **DFadmin** and **DFATK**.

Message	Cause	Solution
The DFdiscover Server is not reachable. Confirm the server name and network connection.	There is no internet connection available or server address/name entered is invalid.	Confirm that your internet service is functioning properly. Confirm that you have correctly typed DFdiscover Server in the login dialog. Otherwise, contact your DFdiscover Study Coordinator or System Administrator.
The DFdiscover Server is not reachable. Confirm the server name and network connection.	There is no internet connection available or server address/name entered is invalid.	Confirm that your internet service is functioning properly. Confirm that you have correctly typed DFdiscover Server in the login dialog. Otherwise, contact your DFdiscover Study Coordinator or System Administrator.
The connection to the server was closed! Please re-login.	Possible Cause: While using the Password Reset dialog, the internet connection was lost.	Confirm that your internet connection is functioning properly, then exit and try again. Otherwise, contact your DFdiscover Study Coordinator or System Administrator.
Unable to connect to DFdiscover server! - Check if DFdiscover EDC Service is running!	Possible Cause: DFdiscover is not running on the specified server or DFdiscover EDC Service is not running.	Contact your DFdiscover Study Coordinator or System Administrator and confirm that the DFdiscover Server and DFdiscover EDC Service is running properly.
Error: No DFdiscover connection available - contact administrator	DFdiscover is not running on the specified server or DFdiscover EDC Service is not running.	Contact your DFdiscover Study Coordinator or System Administrator and confirm that the DFdiscover Server and DFdiscover EDC Service is running properly.
Authorization Failed - incorrect username or password	Username or Password were either mistyped or could not be validated by the server.	Re-enter your Login and Password, and attempt to login again. Contact your DFdiscover Study Coordinator or System Administrator and confirm that your login account is active.
Account locked because of repeated password errors...	Possible Cause: Multiple attempts to login with incorrect password have blocked the account.	Contact your DFdiscover Study Coordinator or System Administrator and request re-activation of your account.
Error: Not Administrator	Attempting to login to DFadmin application using a non-admin account.	Only Study and System Administrators can use the
DFadmin application. Contact your DFdiscover System Administrator.		

Message	Cause	Solution
... - Check if DFdiscover server is running!	Connected to DFdiscover EDC service however DFdiscover Server is not available.	Contact your DFdiscover System Administrator. Confirm that the DFdiscover Server and DFdiscover EDC Service is running properly.
Error: This client software is not compatible with this server	Connecting to DFdiscover server with incompatible version of DFdiscover client application.	Contact your DFdiscover System Administrator to download a version of the DFdiscover application that is compatible with the DFdiscover server.
Error: No DFdiscover license available	All licensed seats are in use by other DFdiscover users.	Wait for a seat to become available and try again. If the problem persists or recurs frequently contact your Study or System Administrator to suggest they consider adding additional seats to the DFdiscover server.

Locks

DFdiscover creates a lock directory for those operations or activities that require exclusive access to a resource. **DFdiscover** creates the lock directory at the appropriate time and then deletes the lock directory when exclusive access is no longer required. In unusual circumstances, it may occur that the process that is controlling the lock directory exits prematurely, leaving the lock directory behind. In these cases, it will be necessary to remove the lock directory.

```
/opt/dfdiscover/work/DFmaster.rpcd.lock
```

It is essential to the correct operation of **DFdiscover** that only one copy of the **DFmaster.rpcd** process be running at any one time. When a **DFmaster.rpcd** process is already running, the **DFdiscover** system and the UNIX system (through the **portmapper**) are able to enforce this requirement. However, when **DFmaster.rpcd** is starting and before it has a chance to register itself with the portmapper, it may occur that two independent processes or users attempt to start more than one **DFmaster.rpcd**. To prevent this from occurring, the **DFmaster.rpcd** process tests for and then creates the lock directory, `/opt/dfdiscover/work/DFmaster.rpcd.lock`, at the beginning of its initialization. If the lock directory already exists, the following message appears:

```
DFmaster.rpcd[25859]: another process is starting a master daemon or the
daemon has been disabled - waiting 20 seconds before retry...
DFmaster.rpcd[25859]: another process is starting a master daemon or the daemon
has been disabled - waiting 40 seconds before retry...
DFmaster.rpcd[25859]: another process is starting a master daemon or the daemon
has been disabled - waiting 60 seconds before retry...
DFmaster.rpcd[25859]: another process is starting a master daemon or the daemon
has been disabled - waiting 80 seconds before retry...
DFmaster.rpcd: failed to start service
```

and the **DFmaster.rpcd** process exits without starting (as another **DFmaster.rpcd** is presumably starting at the same time). If the lock directory does not exist, it is created and subsequently removed by **DFmaster.rpcd** after it has successfully registered the UNIX portmapper.

```
### /tmp/DFslave.rpcd.hostname.lock
```

When a slave process starts, it follows the same sequence of initialization steps that the master does, and it too needs to ensure that only one slave process is executing. In the case of the slave process it is however true that there can be multiple slaves across a **DFdiscover** network, but never more than one on any given computer. Hence, a starting slave creates a lock directory in a directory that is local to the machine rather than in shared directory space. In particular, the lock is created as `/tmp/DFslave.rpcd.hostname.lock`.

On (mis-configured) systems where `/tmp` is not writable by everyone, a **DFdiscover** slave cannot be started (there will likely be other non-**DFdiscover** problems as well). This is an administrative issue to keep in mind if the `/tmp` directory is modified.

Diagnosing a study server

Occasionally, as a result of unusual circumstances, a study server may exit and subsequently fail to restart. In such a case, the **DFstudyDiag** utility is used to determine why the study server will not start. The utility begins a series of checks for consistency of server status. The results of the checks are reported and may require root privileges to complete the diagnosis.

It is safe to execute this utility at any time, however, in a larger networked environment with many slave machines, the diagnosis may require a long period of time to complete.

Diagnosing a study server

1. Execute **DFstudyDiag**

DFstudyDiag requires the study number of the study to diagnose. To start the utility, type:

```
# /opt/dfdiscover/utls/DFstudyDiag -s ###
```

where ### is the study number.

2. Monitor the output from the diagnosis checks

It is possible that the diagnosis will find that the study server is fully operational, as in:

```
# /opt/dfdiscover/utis/DFstudyDiag -s 253
Diagnosing study server 253 starting Mon Aug 14 09:36:09 2017...

>> Trying to contact study server directly...
<< Study server is currently operational and responding.
```

3. Apply the suggested corrective action

If the diagnosis detects a problem, it suggests a solution, as in this example:

```
# /opt/dfdiscover/utis/DFstudyDiag -s 253
Diagnosing study server 253 starting Mon Aug 07 09:41:27 2017...

>> Trying to contact study server directly...
<< Failed.

>> Trying to load studies database from master...
<< OK.

>> Contacting slaves on candidate hosts...

>> Checking portmapper entries on candidate hosts...
<< OK.

>> Looking for existing serverstatus file...
<< The file '/opt/dfdiscover/work/.serverstatus253' exists although no study server
<< appears to be running. The file should be removed.
```

Please show this output to your DFdiscover administrator.

Implement the suggested solution, and then attempt diagnosis again. Some solutions may require multiple steps that will not be detected by only a single diagnosis.

Troubleshooting Techniques

Check for **DFdiscover** or system error messages

The **DFdiscover** processes log all unexpected events to the system message log. It is recommended that this file always be checked after an unexpected event. Messages in this file may not always make sense to a user or **DFdiscover** administrator but they can be very valuable to a **DFdiscover** support person.

Permissions: does the problem occur as user datafax?

UNIX permissions are a common source of problems within **DFdiscover**. If a problem occurs with a **DFdiscover** executable, one can easily determine whether or not permissions are a cause of (or contributor to) the problem. As user datafax, re-run the problematic executable, repeating the steps that led to the problem. If the problem is not repeatable, then permissions are likely a cause of the problem.

To verify the **DFdiscover** permissions for a study, use the **DFstudyPerms** utility. This application has the following usage:

```
% /opt/dfdiscover/utis/DFstudyPerms [-f] [-g group] #
```

where -f instructs the utility to fix any problems that it finds, and -g group indicates the group name for the study, which is required when the group name is not the default studies.

Permissions: are the **DFdiscover** application permissions correct?

To restore the **DFdiscover** permissions that were applied to the software at the time that it was installed, re-run the **SETPERMS** script that is part of the **DFdiscover** software and can be located in the /opt/dfdiscover directory.

Which step of a shell-script is failing?

All of the **DFdiscover** shell scripts are relatively quiet during their execution. They echo the occasional message. To make the shell script very verbose, add the statement:

```
set -x
```

as the second line of any existing shell script. Re-run the shell script and watch all of the output go by!

Other troubleshooting areas

The remaining material is a compilation of troubleshooting information that has been collected over the years from telephone and web-based support.

Although this section is comprehensive, it is possible that your site may experience a failure or symptom that is not enumerated here. If you are unable to resolve the problem using this resource, contact **DF/Net Research, Inc.** as outlined under [Getting Help](#).

Database server process cannot be started

Attempts to start client applications for a particular study database cause a Study server not running on local network message to appear on the user's screen.

The utility application **DFstudyDiag** must be run with `-s #`, where `#` is the **DFdiscover** study number, to determine the reason that the study server cannot be started. The example below illustrates use of **DFstudyDiag** on study 248.

Using DFstudyDiag to evaluate the status of a study server

```
# /opt/dfdiscover/utis/DFstudyDiag -s 248

Diagnosing study server 248 starting Tue May 8 12:12:25 2018...
>> Trying to contact study server directly...
<< Failed.
>> Trying to load studies database from master...
<< OK.
>> Contacting portmapper on candidate hosts...
<< OK.
>> Contacting slaves on candidate hosts...
<< OK.
>> Checking portmapper entries on candidate hosts...
<< OK.
>> Looking for existing serverstatus file...
<< The file '/opt/dfdiscover/work/.serverstatus248' exists although no study
server
<< appears to be running. The file should be removed.
Please show this output to your DFdiscover administrator.
```

The output will include instructions for resolving the problem.

CRF images are missing, or partially processed documents are found in an incoming daemon work directory

The incoming daemons process incoming documents in their own separate work directories (configured using **DFadmin** - Incoming). When a document has been completely processed the CRF images are moved to either the study CRF pages directory or to the unidentified router. Thus if you find any pages remaining in an incoming daemon work directory after it has exited, something has gone wrong.

This situation may be brought to your attention by a user who indicates that 1 or more pages of a received document appear to be missing or misplaced. They might report that in validating the data records for a new document, they can get the first few pages to come up in **DFexplore**, but the last pages cannot be located. There are messages in the errors file similar to the following:

```
/opt/dfdiscover/work/001/1732: File exists
```

This message is a symptom that a previous **DFinbound.rpc** process terminated before completing the processing of an image.

The **DFinbound.rpc** application starts after the arrival of a new image and exits when processing of the image is complete. During this processing, **DFinbound.rpc** creates and manages several intermediate versions of image files, routing each of them before completion so that the work directory is always empty when it exits (it should also be empty before it starts). If **DFinbound.rpc** is terminated in-progress then it is possible for one or more intermediate files to be left behind. The next time that the **DFinbound.rpc** application processes a document in that working directory it notices that it is not empty and attempts to clean up the remaining files. In general, no action is required (other than to wait until the next image is processed) to recover from this situation.

Recovering from a Power Failure

Your **DFdiscover** server should be paired with an uninterruptible power supply (UPS) that is configured to shut your server down should an interruption in power exceed the capacity of the UPS. If your server is not protected by a UPS or other source of emergency power and power is interrupted, there may be problems restarting your server or gaining access to **DFdiscover**.

The steps needed to recover from a power failure are as follows.

1. Verify that all DFdiscover processes are running

After the system has restarted, check to verify that all expected **DFdiscover** processes are running. This is easily accomplished from the command-line.

```
ps -ef | grep DF
```

The following output is typical of a running **DFdiscover** server (*long output lines have been truncated*).

```
datafax 10369 5190 0 06:30:09 ?      0:00 /opt/dfdiscover/bin/..x86_linux/bin/DFoutbound.rpc -c ...
datafax 10346 5190 0 06:30:06 ?      0:00 /opt/dfdiscover/bin/..x86_linux/bin/DFmaster.rpcd -c ...
datafax 10370 10369 0 06:30:09 ?      0:00 /opt/dfdiscover/bin/..x86_linux/bin/DFoutbound_B ...
datafax 10356 5190 0 06:30:06 ?      0:00 /opt/dfdiscover/bin/..x86_linux/bin/DFslave.rpcd -q
root 24273 24269 0 13:50:39 pts/5      0:00 grep DF
datafax 10363 5190 0 06:30:09 ?      0:00 /opt/dfdiscover/bin/..x86_linux/bin/DFedcservice
```

2. Set DFuserdb locking to a known state

Use the following command to set the **DFuserdb** locking to a known state. Shut down the **DFdiscover** server and start it up again as follows:

```
DFshutdown
DFuserdb -unlock
DFbootstrap
```

Repeat step 1 to verify that the expected processes are running.

3. Remove any stale study server status files

Before your system lost power, any studies that were open at the time had a study server status file associated with them. These files may need to be removed when power is restored as they are stale and do not reflect the current status of your system. See [Other troubleshooting areas](#) for information diagnosing and fixing study server problems.

4. Remove any stale lock files

Before your system lost power, it may have been in the process of starting up. This can happen in cases where power interruption is intermittent. Review [Locks](#) and remove any stale locks, then restart the **DFdiscover** server using **DFbootstrap**. Repeat step 1 to verify that all expected processes are running.

5. Perform consistency checks on active studies

After your system has started, run all the DF_IC*** reports and fix any problems reported.

Recovering from a full Filesystem

As you have no doubt already learned (or been reminded) in this chapter, disk maintenance in your **DFdiscover** environment is an important and ongoing requirement. To help protect your **DFdiscover** environment from unexpected problems that result from full filesystems, **DFdiscover** shuts down the incoming image process when one or more full filesystems is detected. This does not mean that incoming images are blocked. Images can still be received by the fax modems and **HylaFAX** during this time; it only means that any images received will not be processed through **DFdiscover** until space is made available in the affected filesystems. However, if this condition is left uncorrected for a period of time it can occur that the filesystem containing the incoming directory also fills, and then **HylaFAX** will shut down and incoming transmissions will no longer be received. This particular situation is discussed in greater detail at the end of this section.

Full filesystem problems are typically detected by the **DFdiscover** incoming daemon when attempting to move CRF images from the incoming work directory to a study pages or pages_hd directory. If the partition containing the study pages or pages_hd directory is at 100% capacity or reaches 100% capacity during the processing of an incoming document, **DFdiscover** will fail to complete the processing of the incoming document, it will generate error messages, and finally it will create the directory `/opt/dfdiscover/work/DFfull`. Subsequent to this, each incoming daemon that is started by the arrival of a new document will first check to see if the directory `/opt/dfdiscover/work/DFfull` exists, and if so processing will be aborted. Each new document will remain in the incoming work directory until the situation is corrected.

The steps needed to correct a full filesystem problem are as follows.

1. Determine which filesystem is full

The easiest way to determine which filesystem has reached 100% capacity is to inspect the system error message log file. Scan the messages near the end of the file that have the following appearance:

```
fax /opt/dfdiscover/incoming/fax00123.tif could not be moved to /opt/dfdiscover/archive/1810/0250;
free some disk space and then manually move the file
```

or

```
fax /opt/dfdiscover/work/001/1730/0020002 could not be renamed to /opt/studies/study123/pages/1730/0020002 -
contact your DFdiscover administrator
```

2. Acquire additional disk space in the affected filesystem(s)

If the archive file system is full (as was indicated in the first message), some or all of the existing archive files can be removed to secondary media and then deleted from primary storage. This topic was previously covered in [Archive File Maintenance](#).

If one of the study pages directories is full (as was indicated in the second message), more disk space will need to be acquired.

3. Move any files that the DFdiscover software was unable to move

Once additional free disk space is available, resolve all of the problems that **DFdiscover** noted as errors in the system error message log file. The text of each problem message indicates the step that needs to be performed to resolve the problem. For example, to resolve the problem indicated by the first message, one would:

```
# mv /opt/dfdiscover/incoming/fax00123.tif /opt/dfdiscover/archive/1810/0250
```

and similarly, to resolve the problem indicated by the second message, one would:

```
# mv /opt/dfdiscover/work/001/1730/0020002 /opt/studies/study123/pages/1730/0020002
```

If the filesystem full situation occurred early in the processing of a multi-page document, several of these problems may need to be resolved, typically one per document that could not be moved. Be sure to resolve each of these problems.

4. Remove `/opt/dfdiscover/work/DFfull`

At this point, the problems that may have occurred around the time that a full filesystem condition was detected have been resolved. Removal of the `DFfull` directory will enable **DFdiscover** to resume processing new incoming images as they arrive. Any images that were received during the time that the `DFfull` directory existed are processed automatically on receipt of the next image. To remove the `DFfull` directory:

```
# rmdir /opt/dfdiscover/work/DFfull
```

Interaction with HylaFAX

The above scenario describes what occurs when a **DFdiscover** incoming daemon discovers that there is insufficient disk space available to process an image. In that case, all incoming transmissions are left in the incoming directory defined by your **HylaFAX** configuration. If the problem is left uncorrected for a period of time such that the filesystem containing the incoming directory also fills, then the problem will escalate to the **HylaFAX** level. At the **HylaFAX** level, the **HylaFAX** software will no longer answer any incoming calls. This is now a serious situation as investigators will no longer be able to transmit case report forms to your site.

To correct this problem, you must first resolve the **DFdiscover** filesystem problems that were detected. Follow the steps described above to accomplish this. When these steps are completed the filesystem containing the incoming directory should no longer be at 100% capacity. The next incoming call will be answered by a fax modem and your operation will return to a normal state.


DFdiscover System Administration Tool

DFadmin connects to the **DFdiscover** master via network communication (even if the **DFadmin** is started on the same host as the master) and hence its operation will be affected by network interruptions and/or network failures. The behavior of **DFadmin** in response to such problems is outlined in the following sections:


Auto Logout

There are two situations where auto logout can occur, one expected and one unexpected:

1. Expected - timeout: **DFadmin** will auto logout if there has been no interaction for the amount of time set in the **DFadmin** auto logout dialog, which allows a maximum value of 60 minutes. The following dialog will appear.



2. Unexpected - critical error. Any client application, will auto logout the connected user if a critical error, such as a loss of network connection, is detected. The dialog, **DFexplore** shown here, has the following appearance:



In both instances, the user will be logged out *without* saving any changes. For this reason, users are encouraged to save changes regularly and never leave a client application unattended.

Following an auto logout, the user will be notified on next login that their application auto logged out during its last use and when that occurred.

DFdiscover

If a **DFdiscover** server (master or study database) does not respond to a **DFadmin** request within sixty seconds, a dialog will appear indicating that there has been no response from the **DFdiscover** server. If the suspects that the lengthy delay is a result of network congestion, the user can choose to wait until sixty seconds has elapsed. Otherwise, the user can exit the application (any unsaved changes will be lost).

DFdiscover System Files

This chapter is a reference for the files and directories that comprise a **DFdiscover** installation. It starts with a brief overview of the top level directories of the **DFdiscover** installation directory. It then describes in detail the **DFdiscover** configuration and status files maintained under the *lib* and *work* directories. A similar chapter, [Programmer Guide, DFdiscover Study Files](#), describes the file system maintained for each individual **DFdiscover** study.

DFdiscover system directories

The following directories are part of a standard **DFdiscover** installation.

<i>bin</i>	This directory contains all of the executables included in a DFdiscover installation. All DFdiscover applications are executed from this directory. In most cases, the executable is only a symbolic link to a script that determines and executes the true executable for the current operating system and hardware.
<i>doc</i>	In some cases, the executable is a shell script itself. This directory contains all of the DFdiscover documentation.
<i>ecbin</i>	This directory contains any scripts called by edit checks that are system-wide or common to all studies.
<i>ghostscript</i>	This directory contains configuration information and documentation for the DFgs interpreter included with DFdiscover .
<i>identify</i>	Images received by DFdiscover which cannot be identified as belonging to any of the registered studies, are stored in this directory (1 file per page). They can be viewed, printed, deleted, and if appropriate, routed to a study using the DFexplore Image Router function.
<i>lib</i>	<p>The directory holds all system configuration files (managed using DFadmin), and 3 databases used to keep track of the current status of each study server, the inbound image daemon, and the outbound image daemon. The following files are described in detail later in this chapter in DFdiscover system configuration files</p> <ul style="list-style-type: none"> • DFedcservice.cf EDC service configuration file • DFinbound.cf inbound daemon configuration file • DFinbound.db current status of the inbound daemon • DFlogin.html configuration file for login banner • DFmaster.cf master daemon configuration file • DFoutbound.cf outbound daemon configuration file. • DFoutbound.db current status of the outbound daemon • DFreceipt.sample example auto reply receipt • DFreplyfax.cf configuration file for auto reply feature • DFsites.db database of registered sites for auto reply feature • DFstudies.db current status of all study database servers • DFuserdb.log user, role, and permissions database • DFadmin.db admin sqlite database • iso3166-1 ISO 3166-1 table of 3-character country codes
<i>reports</i>	This directory contains the standard DFdiscover reports which are executable through Reports View in DFexplore .
<i>utils</i>	This directory contains various utility applications. They are less frequently used than the executables in the <i>bin</i> directory. The executable is a symbolic link to a script that determines and executes the true executable for the current operating system and hardware.
<i>work</i>	<p>This directory contains various DFdiscover log and status files. The following files are described in detail later in this chapter in DFdiscover system work files:</p> <ul style="list-style-type: none"> • 001 the work directory used by the incoming image daemon. The directory is used for temporary storage of the incoming pages being processed by the incoming daemon. When image processing is complete this directory will be empty. <p>This directory is not described further below.</p> <ul style="list-style-type: none"> • router_log log of all images in router • fax_log log of all images received • fax_log.idx index on log of all images received • license DFdiscover software license • replyfax_log Log of all receipts sent by auto reply feature • sendfax_log log of all transmitted images • server_log log of server start and stop activity • sitefax_log log of images received from unregistered sites • .serverstatus### Lock for executing server • .seqYYWW incoming image sequence numbering • .seqOUT outgoing image sequence numbering

<code>x86_linux</code>	Operating system and hardware specific executables for Redhat or SUSE Linux on Intel hardware. If DFdiscover is not installed on Linux Intel hardware, this directory will not be present.
------------------------	---

The remainder of this chapter describes the files kept under the **DFdiscover** *lib* and *work* directories. Each file is described using a common format.

WARNING: Do not edit system files.

You should *never* try to manually edit any of the **DFdiscover** system files! These files are created and maintained by one or more **DFdiscover** applications and are integral to the proper functioning of the various applications which together make up the **DFdiscover** system. Manual edits could produce unexpected and disastrous effects. It is however perfectly safe to read these files for system diagnosis, or use them as input for system administration scripts. It is for this reason that this information is provided here.

Format used to describe files

Each file documented in this section is described with the following attributes:

File attributes

Heading	Description
Usual Name	the file name that is usually given to files having this format. Some files are kept at the DFdiscover directory level while others are kept separately with each study directory.
Type	one of: "clear text" or "binary". Clear text files can be reviewed with any text editor.
Created By	the name of the DFdiscover program(s) that create and modify this file. If you need to edit the contents of the file, use the program listed here.
Used By	the name of the DFdiscover program(s) that reference or read this file.
Field Delimiter	how fields within a record are delimited. Typically, the delimiter is a single character.
Record Delimiter	how records within the file are delimited. Typically, the delimiter is a single character.
Comment Delimiter	how comments within the file are delimited. If comments are not permitted within the file, "NA" is indicated.
Fields/Record	the expected number of fields per record. If the number of fields varies across records, the minimum number is given, followed by a \pm .
Description	a detailed description of the meaning of each field.
Example	one or more example records from the file

DFdiscover system configuration files

DFedcservice.cf

Usual Name	<i>DFedcservice.cf</i>
Type	clear text
Created By	INSTALL
Used By	DFedcservice
Field Delimiter	<i>space</i> character
Record Delimiter	*
Comment Delimiter	#
Fields/Record	2

Description	<p>The file configures the DFedcservice processes, telling it where to find its SSL key and certificate as well as which DFmaster.rpcd to talk to. The format of the configuration file is one keyword, value pair per record with one occurrence of each of the following keywords required: [<i>master, key, certificate</i>].</p> <p>The file can also instruct the DFedcservice process to bind to a particular network interface using the optional <i>bindaddr</i> keyword. This is useful if the server running DFedcservice has more than one network adapter (either real or virtual), and it also services non-DFdiscover requests on port 443.</p> <p>The optional <i>emailTempPass</i> config flag can be set to have the system email a user a temporary password rather than a link by default. Add the following line: <code>emailTempPass 1</code>.</p> <p>The optional <i>defaultAPI</i> parameter can be set to direct an API client app to use the specified API server to service API requests. API clients do not need to separately define an API server unless connecting to 5.8 or older APIs.</p> <p>The optional <i>dfengagefqdn</i> parameter is required when DFweb is used to provide DFengage invite and login web links on the ePRO Management page. Add the following line: <code>dfengagefqdn FQDN</code> where FQDN is the fully qualified domain name for the DFengage web instance.</p> <p>The optional <i>azureredirecturi</i> parameter can be set to support Microsoft Azure Entra ID single sign on under specific configurations. This parameter is only required when the DFdiscover server name does not match the DFweb service with the same URL over port 8443.</p> <p>The optional <i>ssologinonly</i> parameter can be set to require all users to login using Microsoft Azure Entra ID single sign on, rather than using DFdiscover credentials. To enable this option, add the following line: <code>ssologinonly 1</code>. To allow login with both Microsoft SSO and DFdiscover DFdiscover credentials, add the following line: <code>ssologinonly 0</code>.</p> <p>A set of parameters is provided for optional configuration of Azure SMS Messaging. These include the following:</p> <ul style="list-style-type: none"> • <i>azuresmsenabled</i>: 1=true or 0=false (default is 0) • <i>azuresmsfrom</i>: sender phone number • <i>azuresmsaccesskey</i>: Azure access key • <i>azuresmsurl</i>: Azure service endpoint • <i>azuresmsmessage</i>: SMS message template including the <code>{{AUTHCODE}}</code> placeholder for the security code <p>Finally, the list of ciphers available for secure communication can be specified using the optional <i>ciphers</i> keyword. The value is a comma-delimited list of ciphers used for secure communication. DFdiscover ships with the current recommended list of ciphers based upon widely-accepted best security practices. The list can be edited as needed to add newer ciphers (should they become available after the DFdiscover release). A good, albeit technical, reference for this can be found here as well as an up-to-date list of recommended ciphers</p>
Example	<pre> master apache.dfnetresearch.com key /opt/dfdiscover/lib/edckey.pem certificate /opt/dfdiscover/lib/edccert.pem bindaddr 192.168.3.30 emailTempPass 1 defaultapi explore.dfdiscover.com:4433 dfengagefqdn dfengage.dfdiscover.com azureredirecturi https://explore.dfdiscover.com:8443/login azuresmsfrom +15551234567 azuresmsaccesskey ABCDEFGHIJKLMNOPQRSTUVWXYZ1234567890 azuresmsurl https://sms.example.azure.com/sms?api-version=2023-03-31 azuresmsmessage Use code {{AUTHCODE}} for DFdiscover verification. \nUtilisez le code {{AUTHCODE}} pour la verification DFdiscover. ciphers ECDHE-RSA-AES128-SHA,ECDHE-RSA-AES128- SHA256,TLS_AES_128_GCM_SHA256,TLS_CHACHA20_POLY1305_SHA256 </pre>

Usual Name	<i>DFinbound.cf</i>
Type	clear text
Created By	DFadmin
Used By	DFinbound.rpc
Field Delimiter	=
Record Delimiter	*
Comment Delimiter	#
Fields/Record	2
Description	The file configures the DFinbound.rpc daemon. The format of the configuration file is one keyword, value pair per record with one occurrence of each of the following keywords required: [<i>WORKING_DIR</i> , <i>INBOUND_ARCHIVE_DIR</i>].
Example	<pre>WORKING_DIR=/opt/dfddiscover/work INBOUND_ARCHIVE_DIR=/opt/dfddiscover/archive</pre>

DFinbound.db

Usual Name	<i>DFinbound.db</i>
Type	clear text
Created By	DFadmin
Used By	DFmaster.rpcd
Field Delimiter	/
Record Delimiter	NA, only one record permitted
Comment Delimiter	#
Fields/Record	7
Description	<p>The incoming daemon database maintains the status of the incoming daemon.</p> <p>The fields within the single record are defined in Field descriptions for <i>DFinbound.db</i></p>
Example	<pre>001 /opt/dfddiscover/bin/DFinbound.rpc -c /opt/dfddiscover/lib/DFinbound.cf purgatory Inbound daemon 001</pre>

Field descriptions for DFinbound.db

Field #	Contains	Description
1	daemon number	the unique identifier of the daemon. The daemon number must be in the range 1 to 255.
2	host name	the host name of the machine that the daemon is executing on. If the daemon is currently processing an image, this field contains the host name of the machine that Dfinbound.rpc is executing on. If the daemon is not running, the field is empty.
3	RPC program number	the RPC program number that the daemon is servicing. If the daemon is not running, the field is empty.
4	RPC version number	the RPC version number that the daemon is servicing. If the daemon is not running, the field is empty.
5	command to start daemon	the command that the slave daemon executes to start the incoming daemon. The leading applications name will always be the same while the argument configuration file may change across daemons.
6	candidate host names	a comma delimited list of machines that the incoming daemon can execute on (no spaces).
7	label	a descriptive label for the daemon. It is generally ignored, and is only provided so that this database format is compatible with the format of the studies database.

DFlogin.html

Usual Name	<i>DFlogin.html</i>
Type	HTML
Created By	DFadmin
Used By	DFedcservice
Field Delimiter	NA
Record Delimiter	NA
Comment Delimiter	NA
Fields/Record	NA
Description	<p>The login banner HTML file defines the information that is displayed at the bottom of the login screen for DFexplore, DFsetup, DFadmin and DFsend. This file may contain only HTML text as configured in DFadmin.</p> <p>For more information, refer to Configuring the DFdiscover master.</p>
Example	<pre><html> <center> <h1>DF/Net Research, Inc.</h1> <p><h3>Seattle, Washington, USA</h3></p> <p>For assistance, contact help@dfnetresearch.com<p> </center> </html></pre>

DFmaster.cf

Usual Name	<i>DFmaster.cf</i>
Type	clear text
Created By	DFadmin
Used By	DFmaster.rpcd
Field Delimiter	=
Record Delimiter	*
Comment Delimiter	#
Fields/Record	2
Description	<p>This file configures the behavior of the DFmaster.rpcd process. The format of the configuration file is one keyword, value pair per record with one occurrence of each of the following keywords required: [<i>AUTO_SLAVES</i>, <i>CDN_URL</i>, <i>MAILEE</i>, <i>PASSWORD_RULES</i>, <i>PRINTER</i>, <i>VERSION_STRICT</i>].</p> <p>The value for <i>PASSWORD_RULES</i> is an 8-tuple where fields have the following meaning: <i>length, complexity, expiry, reuse, lockout, failures, email, reset</i></p> <ul style="list-style-type: none"> • length the minimum number of characters required in a valid password • complexity the sum of the following complexity rule values: <ul style="list-style-type: none"> ◦ 1 at least one character and one digit must be included ◦ 2 at least one lowercase and one uppercase character must be included ◦ 4 at least one special character must be included • expiry a single password is valid for this number of days and must be changed before this period ends, otherwise the password expires • reuse reuse of this many most recent passwords is not permitted • lockout after multiple failed login attempts, lockout the user for this many minutes • failures the user is locked out if there are this many consecutive failed login attempts • email if the user is locked out, send a notification email to this address • reset allow the user to reset their own password (1) from the login dialog or disallow (0) this functionality <p>Additional requirements for the values matching each keyword are given in Configuring the DFdiscover master.</p>
Example	<pre>AUTO_SLAVES=purgatory MAILEE=root PRINTER=hp5lj CDN_URL=https://cdn.dfdiscover.com/v8 PASSWORD_RULES=8,1,90,3,5,5,user@domain.com,1 VERSION_STRICT=0</pre>

DFoutbound.cf

Usual Name	<i>DFoutbound.cf</i>
Type	clear text
Created By	DFadmin
Used By	DFoutbound.rpc
Field Delimiter	=
Record Delimiter	*
Comment Delimiter	#
Fields/Record	4
Description	<p>The file configures the DFoutbound.rpc process. The format of the configuration file is one keyword, value pair per record with one occurrence of each of the following keywords required: [<i>WORKING_DIR</i>, <i>POLLING_INTERVAL</i>, <i>ROUND_TRIP_TIMEOUT</i>].</p>
Example	<pre>WORKING_DIR=/opt/dfdiscover/work POLLING_INTERVAL=20 ROUND_TRIP_TIMEOUT=24</pre>

DFoutbound.db

Usual Name	<i>DFoutbound.db</i>
Type	clear text
Created By	DFadmin
Used By	DFmaster.rpcd
Field Delimiter	/
Record Delimiter	NA, only one record permitted
Comment Delimiter	#
Fields/Record	7
Description	The outgoing daemon database maintains the status of the outgoing daemon. The fields within the record are defined in Field Descriptions for DFoutbound.db
Example	001 /opt/dfdiscover/bin/DFoutbound.rpc -c /opt/dfdiscover/lib/DFoutbound.cf purgatory Outbound daemon 001

Field descriptions for DFoutbound.db

Field #	Contains	Description
1	daemon number	the unique identifier of the daemon. The daemon number must be in the range 1 to 255.
2	host name	the host name of the machine that the daemon is executing on. If the daemon is currently running, this field contains the host name of the machine that DFoutbound.rpc is executing on. If the daemon is not running, the field is empty.
3	RPC program number	the RPC program number that the daemon is servicing. If the daemon is not running, the field is empty.
4	RPC version number	the RPC version number that the daemon is servicing. If the daemon is not running, the field is empty.
5	command to start daemon	the command that the slave daemon executes to start the outgoing daemon. The leading applications name will always be the same while the argument configuration file may change across daemons.
6	candidate host names	a comma delimited list of machines that the outgoing daemon can execute on (no spaces).
7	label	a descriptive label for the daemon. It is generally ignored, and is only provided so that this database format is compatible with the format of the studies database.

DFprinters

Usual Name	<i>DFprinters</i>
Type	clear text
Created By	text editor
Used By	DFexplore
Field Delimiter	NA
Record Delimiter	*
Comment Delimiter	NA
Fields/Record	1
Description	<p>This optional file defines the list of printer names that are accessible from DFdiscover client or server machines and are PostScript capable. The file contains one printer name per record (line). In each application that has a print dialog, the list of available printers is created from the contents of this file.</p> <p>If the file does not exist, or is empty, the default behavior is used which involves querying the list of printers known to the host via the [lpstat -t] command.</p>
Example	<pre>hp4000_med hplj5_safety hp4050</pre>

DFreceipt.sample

Usual Name	<i>DFreceipt.sample</i>
Type	clear text
Created By	text editor
Used By	DFinbound.rpc
Field Delimiter	NA
Record Delimiter	NA
Comment Delimiter	NA
Fields/Record	NA
Description	<p>A receipt file contains a text message that will be sent back to those clinical sites that are defined to use DFdiscover's Auto Reply feature.</p> <p>A receipt file may contain only ASCII text. Certain substitutions can be performed by the auto reply feature in the body of the message before the receipt is sent. These substitutions are:</p> <ul style="list-style-type: none"> • %s site name as taken from field 2 of <i>DFsites.db</i> • %p total number of pages received in the document • %t local time of day that the document was received • %d local date that the document was received • %g the G3 fax identifier, typically the subscriber id of the sending document
Example	<pre>DFDISCOVER RECEIPT for STUDY XYZ To: %s From: XYZ STUDY coordinating site Re: recently received document from number %g A %p page document was received at %t on %d (local time). If this does not agree with your submission log, contact: John Doe at 1-888-123-4567, extension 7777 THANK YOU FOR YOUR ONGOING STUDY PARTICIPATION.</pre>

DFreplyfax.cf

Usual Name	<i>DFreplyfax.cf</i>
Type	clear text
Created By	DFadmin
Used By	DFinbound.rpc
Field Delimiter	<i>space</i> character
Record Delimiter	*
Comment Delimiter	NA
Fields/Record	2+
Description	This file records the settings from the Auto Reply dialog in DFadmin . The structure of the file is one setting per record where each setting has a keyword to identify it followed by 1 or more arguments. The keywords and their possible arguments are listed in the table "Field descriptions for <i>DFreplyfax.cf</i> below.
Example	<p>The following Auto Reply configuration file has receipts enabled, 3 retries at 5 minute intervals, logging of auto reply failures only, notification to root@servername.com of auto reply failures, processing of all images, and logging of images from unregistered sites only.</p> <pre> receipts 1 attempts 3 5 log 2 notify 1 root@servername.com processing 0 unknownfax 1 0 </pre>

Field descriptions for *DFreplyfax.cf*

Keyword	Argument 1	Argument 2	Argument 3
receipts	0 (disable auto reply feature) or 1 (enable auto reply feature)		
attempts	# (number of retries - default is 1)	# (minutes between retries - default is 10)	
log	0 (do not log auto replies) or 1 (log all auto replies) or 2 (log only auto reply failures)		
notify	0 (do not notify) or 1 (notify)	email address to send notification to if first argument is 1	
processing	0 (process all incoming images) or 1 (process images from registered sites only)		
unknownfax	0 (do not log images from unregistered sites) or 1 (log images from unregistered sites)	0 (do not send a notification email) or 1 (send a notification email)	email address to send notification of unregistered images to if second argument is 1

DFsites.db

Usual Name	<i>DFsites.db</i>
Type	clear text
Created By	text editor
Used By	DFinbound.rpc
Field Delimiter	/
Record Delimiter	*
Comment Delimiter	NA
Fields/Record	5
Description	This file identifies the clinical sites that are to receive receipts when the Auto Reply feature of DFdiscover is enabled. Sites that do not require receipts do not need to be defined in this file. The file contains 1 record per clinical site as described in Field descriptions for DFsites.db
Example	<p>The following example database illustrates two sites where each site receives a unique receipt. The first site (General Hospital) has a fax number as its sender ID and receives a receipt via fax. The second site (St. Elsewhere) has an email address as its sender ID and receives a receipt via email.</p> <pre>1 General Hospital 888-123-4567 1-888-123-4567 opt/studies/xyz/lib/reply.001 1 St. Elsewhere elsewhere@hospital.co mailto:person@hospital.com opt/studies/abc/lib/reply.001</pre>

Field descriptions for DFsites.db

Field #	Contains	Description
1	status	the status field is set to 1 to enable receipts to this site. To disable receipts (but keep the site in the database), set the status field to 0.
2	site name	the name of the site.
3	sender ID	the sender ID (usually the site's email address or fax number) must be included in each received email/fax from the site. It is required to uniquely identify where each transmission came from. Any emails or faxes from the site which do not include this identifier will not be recognized as arriving from this site. This field can accommodate up to a maximum of 100 characters. Comparison of the sender identification with this value is always performed case-insensitive.
4	reply fax number	the email address or fax number to which all auto reply receipts are to be transmitted for this site. It is valid for this email address or fax number to be different from the sender id. Exactly one reply address is permitted in this field, up to a maximum of 100 characters in length. The number of reply addresses is not enforced, therefore the user must exercise care when entering a value for this field.
	receipt file name	the full pathname of the file that contains the auto reply receipt for this site.

DFstudies.db

Usual Name	<i>DFstudies.db</i>
Type	clear text
Created By	DFadmin
Used By	DFmaster.rpcd
Field Delimiter	/
Record Delimiter	*
Comment Delimiter	#
Fields/Record	7
Description	The study servers database maintains the status of all of the study servers defined for your DFdiscover site. The fields within each record are defined as described in Field descriptions for DFstudies.db
Example	102 goofy 536871270 1 /opt/dfdiscover/bin/DFserver.rpc -c /local/study102/lib/DFserver.cf goofy,ariel Example 102

Field descriptions for DFstudies.db

Field #	Contains	Description
1	daemon number	the study number of the server. The study number must be in the range 1 to 999. Study numbers beyond 255 are available for EDC studies and only if the software is licensed to allow it.
2	host name	the host name of the machine that the server is executing on. If the server is currently running, this field contains the host name of the machine that DFserver.rpc is executing on. If the server is not running, the field is empty.
3	RPC program number	the RPC program number that the server is servicing. If the server is not running, the field is empty.
4	RPC version number	the RPC version number that the server is servicing. If the server is not running, the field is empty.
5	command to start server	the command that the slave daemon executes to start the study server. The leading applications name will always be the same while the argument configuration file will change across study servers.
6	candidate host names	a comma delimited list of machines that the study server can execute on (no spaces).
7	label	the study name

DFstudyspaces.db

Usual Name	<i>DFstudyspaces.db</i>
Type	clear text
Created By	DFserveradmin
Used By	DFadmin, DFmigrate
Field Delimiter	/
Record Delimiter	*
Comment Delimiter	#
Fields/Record	3
Description	<p>This file holds the UNIX pathnames where users are allowed to install DFdiscover study directories.</p> <p>Each record is comprised of 3 fields. The first field contains a label displayed in DFadmin when users select the 'Study Space' for a new study, the second field contains the full UNIX pathname corresponding to that label, and the third is either a 0 to indicate this study space is not available for new studies, or a 1 if it is. The UNIX pathname is the root directory for all DFdiscover studies created under that space.</p>
Example	<p>NIH Studies/opt/studies1 1 Other Studies/opt/studies2 1</p> <p>To create a new study a user selects a study space and then enters a study folder name. For example, selecting 'NIH Studies' and then entering ECIC as the folder name would create study directory /opt/studies1/ECIC.</p>

DFuserdb.log

Field #	Contains	Description
1	Record Time Stamp	the date and time this record was created
2	Record Type	the static 4 letter record identifier: USER
3	Record Modifier	the username of the user who created this record
4	User Login (key field)	the username uniquely identifying a DFdiscover user even if that user has been deleted or made inactive
5	Status	the status of this username (0=Deleted, 1=Inactive or 2=Active)
6	Full Name	full name (Maximum 256 characters)
7	Affiliation	affiliation name (Maximum 256 characters)
8	Address	street address (Maximum 256 characters)
9	City	city name (Maximum 256 characters)
10	State	state or province (Maximum 256 characters)
11	Postal Code	ZIP or postal code (Maximum 256 characters)
12	Country	country (Maximum 256 characters)
13	Telephone	phone number (Maximum 256 characters)
14	Fax	fax number (Maximum 256 characters)
15	Email	email address (Maximum 256 characters)
16	Language	language preference (0=English or 1=French) <i>[[ignored in the current release]]</i>
17	Receipt	how user prefers to receive messages (0=Mail, 1=Fax or 2=Email) <i>[[ignored in the current release]]</i>
18	Router Access	permission to use DFrouter (0=No or 1=Yes)
19	Administrator	permission to perform administrative tasks at the system or study levels (blank=No Admin privileges, -1= DFdiscover Admin, *=Admin for All Studies, or list of study numbers=Admin for specified studies)

Field descriptions for *DFuserdb.log* - User Role Record

Field #	Contains	Description
1	Record Time Stamp	the date and time this record was created
2	Record Type	the static 4 letter record identifier: USRL
3	Record Modifier	the username of the user who created this record
4	User Login (key field)	the username uniquely identifying a DFdiscover user even if that user has been deleted or made inactive
5	Role ID (key field)	a sequential identification number, unique across all roles for all studies, assigned to each role when it is created, regardless of whether or not the role is saved
6	Instance (key field)	a sequential number which identifies the potentially multiple specifications of the current study role for the current user (i.e. for different sites and/or subjects) on different rows in the user permissions table
7	Status	the status of this instance of this role for this user (0=Deleted, 1=Inactive or 2=Active)
8	Sites	sites at which this user has this role as specified by this instance
9	Subjects	subjects for which this user has this role as specified by this instance

Field descriptions for *DFuserdb.log* - User ePRO Admin Permission Record

Field #	Contains	Description
1	Record Time Stamp	the date and time this record was created
2	Record Type	the static 4 letter record identifier: EPRM
3	Record Modifier	the username of the user who created this record
4	User Login (key field)	the username uniquely identifying a DFdiscover user even if that user has been deleted or made inactive
5	Study	a DFdiscover study number
6	Status	the status of this permission for this user (0=Deleted, 1=Inactive or 2=Active)
7	Sites	sites at which this user has this role as specified by this instance
8	User Type	the type of ePRO admin permission assigned: 1=Admin, 2=Monitor or 3=Viewer

Field descriptions for *DFuserdb.log* - Role Record

Field #	Contains	Description
1	Record Time Stamp	the date and time this record was created
2	Record Type	the static letter record identifier: ROLE

Field #	Contains	Description
3	Record Modifier	the username of the user who created this record
4	Role ID (key field)	a sequential identification number, unique across all roles for all studies, assigned to each role when it is created, regardless of whether or not the role is saved
5	Study	a DFdiscover study number
6	Status	the status of this role (0=Deleted, 1=Inactive or 2=Active)
7	Role Name	the role name, which must be unique within a study, but not across studies (Maximum 256 characters)
8	Description	the role description (Maximum 256 characters)
9	DFexplore Views	the list of Explore Views accessible by users with this role (0=None, 1=Views-Info, 2=Views-Schedule, 3=Views-Data, 4=Views-Queries, 5=Views-History, 6=Views-Reasons, 7=Views-Image, 8=Views-Data - with Select, 9=Views-Reports, 10=Views-Status, 11=Views-List, 12=SavePrint-Blank CRFs, 13=SavePrint-Completed CRFs, 14=SavePrint-Images, 15=SavePrint-Data, 16=SavePrint-Reports, 17=SavePrint-DFsas jobs, 18=Misc-Developer, 19=Reports-Create lists, 20=List-Create views, 21=Image-Delete new records, 22=Image-Raw data entry, 23=May disable edit checks, 24=List-Import Data, 25=Image-Create tasks, 26=Show query status detail, 27=Data-Import subject CRFs, 28=Data-Submit PDF, 29=Dashboard, 30=Data-Attach subject document or * represents ALL views)
10	Reports	the list of reports accessible by users with this role (DFdiscover report names in comma delimited format or * represents ALL reports)
11	Study Reports	the list of study reports accessible by users with this role (study report names in comma delimited format or * represents ALL study reports)
12	Tasks	the list of Tasks accessible by users with this role (study task names in comma delimited format or * represents ALL tasks) <i>[not implemented in this release]</i>

Field #	Contains	Description
13	Tools	the comma-delimited list of DFdiscover applications (or application components) accessible by users with this role (0=None, 2=Setup-View, 8=Server-Export Data, 9=Server-Batch Edit checks, 11=Server-Import Data, 12=Setup-Plates, 13=Setup-Edit checks, 14=Setup-Lookup Tables, 15=Setup-Centers, 16=Setup-Missing Map, 17=Setup-Sort Map, 18=Setup-Visit Map, 19=Setup-Conditional Terminations, 20=Setup-Conditional Cycles, 21=Setup-Conditional Visits, 22=Setup-Conditional Plates, 23=Setup-Query Titles, 24=Setup-Query Covers, 25=Setup-Query Messages, 26=Setup-CRF Type Map, 27=Setup-CRF Background Map, 28=Setup-Query Category Map or * represents ALL DFdiscover applications)
14	Auto Logout Default	the initial default auto logout period (in minutes) for DFexplore and DFsetup users with this role (The 'Maximum' can be any value from 1 to 1440 minutes). (0 represents inheriting this value from Study settings).
15	Auto Logout Maximum	the maximum value that can be set as auto logout period (in minutes) in DFexplore and DFsetup by users with this role (The 'Maximum' can be any value from 1 to 1440 minutes). (0 represents inheriting this value from Study settings).

Field descriptions for *DFuserdb.log* - Role Perm Record

Field #	Contains	Description
1	Record Time Stamp	the date and time this record was created
2	Record Type	the static 4 letter record identifier: RLPM
3	Record Modifier	the username of the user who created this record
4	Role ID (key field)	a sequential identification number, unique across all roles for all studies, assigned to each role when it is created, regardless of whether or not the role is saved
5	Instance (key field)	a sequential number which identifies the different specifications making up the definition of the current role, i.e. the different rows in the role permissions table for the same Role ID
6	Status	the status of this role instance specification (0=Deleted, 1=Inactive or 2=Active)
7	Visits	the visits to which these permissions apply
8	Plates	the plates to which these permissions apply

Field #	Contains	Description
9	Data Permissions	a list of permissions for data records (1=Create, 2=Modify, 3=Delete, 4=Register as Lost) which are applied in DFexplore and DFweb .
10	Query Permissions	a list of permissions for queries which are applied in DFexplore and DFweb . Operations allowed both manually and by edit checks: (1=Create, 2=Modify, 3=Delete, 4=Approve replies, 5=Create/Modify replies), operations allowed only when performed by edit checks: (6=Create, 7=Modify, 8=Delete, 9=Approve replies, 10=Create/Modify replies)
11	Reason Permissions	a list of permissions for reason records which are applied in DFexplore and DFweb . Operations allowed both manually and by edit checks: (1=Create, 2=Modify, 3=Delete, 4=Approve), operations allowed only when performed by edit checks: (6=Create, 7=Modify, 8=Delete, 9=Approve)
12	Get Levels	only records at these levels can be retrieved in the data collection tools. (* represents ALL levels, equivalent to 0-7)
13	Modify Levels	only records at these levels can be modified in the data collection tools. (* represents ALL levels, equivalent to 0-7)
14	Write Levels	records can be written to only these levels in the data collection tools. (* represents ALL levels, equivalent to 0-7)
15	Hidden Fields	permission to see Hidden Fields - applied in the data collection tools. (0=None or 1=All)
16	Internal Queries	permission to see Internal queries - applied in the data collection tools. (0=None, 1=Resolved queries only or 2=All)

Field descriptions for *DFuserdb.log* - Role Restrictions Record

Field #	Contains	Description
1	Record Time Stamp	the date and time this record was created
2	Record Type	the static 4 letter record identifier: QCAT
3	Record Modifier	the username of the user who created this record
4	Study (key field)	a DFdiscover study number
5	Role ID (key field)	a sequential identification number, unique across all roles for all studies, assigned to each role when it is created, regardless of whether or not the role is saved
6	Status	the status of this role instance specification (0=Deleted, 1=Inactive or 2=Active)
7	Visits	the visits to which these restrictions apply
8	Plates	the plates to which these restrictions apply
9	No Create	a list of query category codes which may not be created in DFexplore , DFweb , and DFcollect
10	No Modify	a list of query category codes which may not be modified in DFexplore , DFweb , and DFcollect
11	No Delete	a list of query category codes which may not be deleted in DFexplore , DFweb , and DFcollect
12	No Approve	a list of query category codes which may not be approved in DFexplore , DFweb , and DFcollect
13	No Reply	a list of query category codes which may not be replied to in DFexplore , DFweb , and DFcollect

Field descriptions for *DFuserdb.log* - Password Record

Field #	Contains	Description
1	Record Time Stamp	the date and time this record was created
2	Record Type	the static letter record identifier: PASS
3	Record Modifier	the username of the user who created this record
4	User Login (key field)	the username uniquely identifying a DFdiscover user even if that user has been deleted or made inactive.
5	Password Hash	the MD5 hash of the user login and password
6	Expiry	the time in seconds since Jan 1, 1970 that this password expires.

DFAdmin.db

Usual Name	DFadmin.db
Type	binary
Created By	DFserver.rpc, DFadmindb
Used By	DFserver.rpc, DFedcservice, DFadmin
Field Delimiter	NA
Record Delimiter	NA
Comment Delimiter	NA
Fields/Record	NA
Description	<p>All user and role information and change history is stored in a sqlite database. The database contains the following tables:</p> <ul style="list-style-type: none"> • DFlogin: user profiles • DFdeviceLog: devices for 2FA • DFeproAdmin: user ePRO admin permissions • DFeproUser: ePRO users • DFeproUserLog: ePRO user logs • DFuserStudy: users assigned to study roles • DFroleStudy: roles defined in studies • DFqcRestrict: query category restrictions by role • DFpermLog: change history • DFactivityLog: activity by client program (currently DFengage activity only) <p>DFserver.rpc updates DFadmin.db when users and roles are updated.</p> <p>DFedcservice and DFadmin user DFadmin.db to display user and role history.</p> <p>DFadmindb utility converts existing user and role audit trail to DFadmin.db.</p>

iso3166-1

Usual Name	iso3166-1
Type	clear text
Created By	NA
Used By	DFsetup, DFexplore
Field Delimiter	/
Record Delimiter	*
Comment Delimiter	NA
Fields/Record	2
Description	<p>This table contains the official ISO 3166-1 three-letter country codes and country names. It is a static table, read only by DFsetup (when identifying the country for a site) and by DFexplore and DFweb (when reporting data grouped by country).</p>

DFdiscover system work files

router_log

Usual Name	<i>router_log</i>
Type	clear text
Created By	DFmaster.rpcd
Used By	DFmaster.rpcd, DFserver.rpc
Field Delimiter	
Record Delimiter	*
Comment Delimiter	NA
Fields/Record	9
Description	<p>DFexplore Image Router activity, including both input and output, is recorded in the <i>*router_log*</i>. Each action performed is appended to the log file. The fields of the log file are described in Field descriptions for router_log</p> <p>The fifth field records the type of operation being performed as follows:</p> <p>cut in two: Indicates that this page was cut resulting in two separate pages.</p> <p>cut from ####.###X###: This page was created from the cut operation performed on the specified page.</p> <p>to study: The page has been sent to the new queue of the study specified in the sixth field.</p> <p>deleted: This page has been permanently deleted.</p> <p>printed: This page was either printed or exported as PDF.</p> <p>from incoming: This page was sent to the router because [DFinbound.rpc] daemon could not identify it (Barcode read failure).</p> <p>Moved from study ###: From DFexplore the page has been sent to the router. It did not belong to the original study.</p> <p>The last four fields of the log record are not relevant for some operations.</p>
Examples	<pre>2-05 10:56:54 sandra 1749.000D001 to study 253 2603 1 3 7-12-07 11:00:39 sandra 1749.000A001 printed 7-12-07 16:33:01 sandra 1749.000A001 deleted </pre>

Field descriptions for router_log

Field #	Description
1	date the operation was performed
2	time the operation was performed
3	user who performed the operation
4	image id on which the operation was performed
5	description of the operation
6	study number
7	subject ID
8	visit number
9	plate number

fax_log

Usual Name	<i>fax_log</i>
Type	clear text
Created By	DFmaster.rpcd
Used By	DF_XXtime, DF_ATfaxes, DF_WFcrfsperwk, DFexplore
Field Delimiter	/
Record Delimiter	*
Comment Delimiter	NA
Fields/Record	5
Description	<p>DFmaster.rpcd appends a record to <i>fax_log</i> for every image received by DFdiscover, thus the most recent transmissions appear at the end of this file.</p> <p>Each record contains 5 pieces of information described in Field descriptions for *fax_log</p> <p>The third field recorded in <i>fax_log</i> is the date and time the transmission was received. For a typical image transmission the date and time comes from the TIFF file "Date & Time" tag created by HylaFAX when the image is received to disk. If DFdiscover receives a TIFF or PDF file as an email attachment, the date and time written to <i>fax_log</i> comes from the computer clock on the machine that processed the incoming mail. This is true, even if a TIFF file attachment already contains a date and time tag representing when it was created at the site. In such cases the original date and time tag is replaced in the TIFF file before it is saved in the DFdiscover archive directory. In the case of a PDF file attachment, Ghostscript is used to convert the file to TIFF format, with the date and time tag set from the machine that processed the incoming PDF, before the attachment is saved in the DFdiscover archive directory.</p> <p>The fourth field contains a quoted string which identifies the sender. It will contain one of the following:</p> <ul style="list-style-type: none"> • the sender id programmed into the sending fax machine if the pages were sent by fax • a blank string if the sender id is unknown because the sender id was not programmed into the sending fax machine • the senders email address if the file was sent as an email attachment, unless the file contains a non-blank image description field (e.g. TIFF files sent by email may contain such information), in which case the image description is used • DFexplore Submit PDF:username if the pages were transmitted using the DFexplore 'Submit PDF' feature.
Example	<code>/opt/dfdiscover/archive/1742/0024 16 Tue Oct 17 04:46:23 2017 555-1212 28</code>

Field descriptions for fax_log

Field #	Description
1	full pathname of image file, typically in the archive directory
2	total number of pages in document
3	date and time document was transmitted
4	sender identification or a combination of <i>app name method:username:pathname</i> (compacted if the pathname exceeds approximately 75 characters or bytes)
5	archive duration. This field is currently ignored unless its value is 0

fax_log.idx

Usual Name	<i>fax_log.idx</i>
Type	binary
Created By	DFmaster.rpcd
Used By	DFmaster.rpcd
Field Delimiter	NA
Record Delimiter	NA
Comment Delimiter	NA
Fields/Record	NA
Description	<p>The <i>fax_log.idx</i> file is an index on the <i>fax_log</i> file and enables DFmaster.rpcd to quickly retrieve individual entries from <i>fax_log</i> for operations such as [Get Context] in DFexplore.</p> <p>This file is not user editable or readable.</p>

license

Usual Name	<i>license</i>
Type	clear text
Created By	DFadmin
Used By	DFmaster.rpcd
Field Delimiter	*tab)
Record Delimiter	*
Comment Delimiter	#
Fields/Record	2
Description	<p>DFdiscover is a licensed software product. It is licensed for a particular machine (by unique host identifier). Each DFdiscover license is typically valid for a one year period. The license is provided yearly by DF/Net Research, Inc. in the form of a codeword. The codeword encodes, among other things, the host identifier, the expiration date, and the maximum number of concurrent users.</p> <p>The license file contains values for each of the following keywords: [<i>hostid</i>, <i>hostname</i>, <i>clients</i>, <i>expiration</i>, <i>password</i>].</p> <p>The value for <i>hostname</i> can be changed without affecting the other values, and in particular, without requiring a new license. However, it is critical that the values for the other keywords not be changed without a new password because they are interrelated.</p>
Example	<pre>hostid SH2X-F77S-JURH-DJMQ-JCHW hostname venus nclients 1 expiration 20061231 password KZTD-CQA3-MY5W-HX5E-24N3</pre>

replyfax_log

Usual Name	<i>replyfax_log</i>
Type	clear text
Created By	DFadmin
Used By	DFoutbound.rpc
Field Delimiter	/
Record Delimiter	*
Comment Delimiter	#
Fields/Record	7
Description	The DFdiscover Auto Reply configuration includes an option to log all receipts sent to registered sites. This log information is stored in this file, one record per receipt, and is described in Field descriptions for replyfax_log
Example	<code>/opt/studies/xyz/lib/DFreceipt.sample SentOk Fri May 18 11:16:23 2018 datafax purgatory 18881234567 1</code>

Field descriptions for replyfax_log

Field #	Description
1	the full pathname of the transmitted receipt
2	the final transmission disposition. Will contain one of two possible values: <i>SentOk</i> or <i>Failed</i>
3	the date and time that the receipt was transmitted
4	the username of the person who requested that the receipt be transmitted
5	the hostname of the machine from which the request was made
6	the email address or fax number that the receipt was sent to
7	the number of attempts needed to successfully send the receipt if the status in field 2 is <i>SentOk</i> , or the total number of failed attempts if the status in field 2 is <i>Failed</i>

sendfax_log

Usual Name	<i>sendfax_log</i>
Type	clear text
Created By	DFoutbound.rpc
Used By	
Field Delimiter	/
Record Delimiter	*
Comment Delimiter	NA
Fields/Record	7
Description	<p>Every document that is sent from DFdiscover is recorded by DFoutbound.rpc to this file. The file is always appended to so that the most recently transmitted document is listed at the end of the file.</p> <p>Each document sent to one recipient represents one record in the file. Each entry records the information described in Field descriptions for sendfax_log</p>
Example	<code>/studyA/reports/QC/001-171203 SentOk Sun Dec 3 11:16:23 2017 root venus 5551212 1</code>

Field descriptions for *sendfax_log

Field #	Description
1	the full pathname of the transmitted file
2	the final transmission disposition. Will contain one of two possible values: <i>SentOk</i> or <i>Failed</i>
3	the date and time of transmission
4	the username of the person who requested the transmission
5	the hostname of the machine from which the request was made
6	the email address or fax number that the transmission was sent to
7	the number of attempts needed to successfully send the transmission if the status in field 2 is <i>SentOk</i> , or the total number of failed attempts if the status in field 2 is <i>Failed</i>

server_log

Usual Name	<i>server_log</i>
Type	clear text
Created By	DFmaster.rpcd
Used By	
Field Delimiter	NA
Record Delimiter	*
Comment Delimiter	NA
Fields/Record	5
Description	<p>Each time that a DFdiscover study server starts or exits, it appends a message to this file in the following format:</p> <p style="padding-left: 40px;">DFserver.rpc.###[PPPPP]: start exit on Host at Date Time</p> <p>where the study number is indicated by ### and the process id by PPPPP. If the server is starting the log message will contain the word <i>start</i> and if the server is exiting, the log message will contain the word <i>exit</i>. The hostname of the CPU that the server started on is reflected in the Host value and the time it started (or exited) in the Date and Time fields.</p> <p>If the server exited normally there should be exactly one pair of <i>start</i> and <i>stop</i> log messages for each unique combination of process id and host.</p> <p>Messages are also appended to this file if a DFdiscover administrator deletes a study, study data, or study data and setup info using the DFadmin 'Delete' option. The example below illustrates the message from each of these operations performed on study 101.</p> <p>This log file is not consulted by any DFdiscover process and is purely for informational purposes. It can be truncated if it grows too large.</p>
Example	<pre>DFserver.rpc.153[20401]: start on idemo41 at Wed Nov 29 22:22:11 2017 DFserver.rpc.153[20401]: exit at Thu Nov 30 00:54:03 2017 DFedcservice.101[21877]: jack@localhost deleted all study data Fri Dec 1 12:00:09 2017 DFedcservice.101[21877]: jack@localhost deleted all study data and setup info Fri Dec 1 12:01:19 2017 DFedcservice.101[21877]: jack@localhost deleted study from datafax registry Fri Dec 1 12:02:58 2017</pre>

sitefax_log

Usual Name	<i>sitefax_log</i>
Type	clear text
Created By	DFoutbound.rpc
Used By	
Field Delimiter	/
Record Delimiter	*
Comment Delimiter	NA
Fields/Record	5
Description	<p>Configuration of the DFdiscover Auto Reply feature allows for logging of each image that is received from an "unregistered" site. An unregistered site is a sending image site that not appear in the <i>DFsites.db</i> database.</p> <p>One log entry is created for each document transmission received from an unregistered site. The format of each entry is described in Field descriptions for sitefax_log</p>
Example	<pre>/opt/dfdcover/archive/1820/0025 1 Fri May 18 10:20:12 2018 "905 522 7284" unregistered</pre>

Field descriptions for sitefax_log

Field	Description
1	full pathname to the location of the received image
2	total number of pages in the received image
3	date and time that the image was received in Day Mon DD HH:MM:SS format
4	Sender identification (typically fax number) of the sending machine
5	always contains the string "unregistered"

.serverstatus###

Usual Name	<i>.serverstatus###</i>
Type	clear text
Created By	DFmaster.rpcd
Used By	DFmaster.rpcd, DFserver.rpc
Field Delimiter	NA
Record Delimiter	NA
Comment Delimiter	NA
Fields/Record	1
Description	<p>This file keeps track of the process ID number of the study server if it is currently executing. It is created and used by DFmaster.rpcd to keep track of which study servers are running.</p> <p>It is a serious error to delete a server status file unless instructed to do so by the output of the DFstudyDiag utility.</p>

.seqYYWW

Usual Name	<i>.seqYYWW</i>
Type	clear text
Created By	DFmaster.rpcd
Used By	DFmaster.rpcd
Field Delimiter	NA
Record Delimiter	NA
Comment Delimiter	NA
Fields/Record	1
Description	<p>This file keeps track of the next sequence number to be assigned to an incoming image within the current week of the current year. These sequence files are created and maintained exclusively by DFmaster.rpcd. Each sequence file has a name of the form <i>.seqYYWW</i> where YYWW represents a particular week within a particular year.</p> <p>Each sequence file contains exactly one positive number, which is the next number to be assigned to the next image which arrives in that particular week of that particular year. When an image arrives, it is assigned the number from the sequence file. The number is then incremented and written back to the file.</p> <p>If the required sequence file does not exist, it is created by the master and should thereafter remain in existence. It is a serious error to delete any sequence file, however if one is accidentally deleted, it can be recreated with the DFcmpSeq utility.</p>

.seqOUT

Usual Name	<i>.seqOUT</i>
Type	clear text
Created By	DFmaster.rpcd
Used By	DFmaster.rpcd
Field Delimiter	NA
Record Delimiter	NA
Comment Delimiter	NA
Fields/Record	1
Description	<p>This file keeps track of the next sequence number to be assigned to the next outgoing image. This sequence file is created and maintained exclusively by DFmaster.rpcd.</p> <p>The sequence file contains exactly one positive number, which is the next number to be assigned to the next image to be sent out. When an image is sent, it is assigned the number from the sequence file. The number is then incremented and written back to the file.</p> <p>If the required sequence file does not exist, it is created by the master and should thereafter remain in existence. It is a serious error to delete any sequence file.</p>

Copyrights

External Software Copyrights

DFdiscover software uses several third-party software components as part of its server side and/or client tools.

The copyright information for each is provided below. If you would like to receive source codes of these third-party components, please send us your request at help@dfnetresearch.com.

DCMTK software package

Copyright© 1994-2011, OFFIS e.V. All rights reserved.

This software and supporting documentation were developed by

OFFIS e.V. R&D Division Health Eschereg 2, 26121 Oldenburg, Germany

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of OFFIS nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Jansson

Copyright© 2009-2014 Petri Lehtinen <petri@digip.org>

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Mimencode

Copyright© 1991 Bell Communications Research, Inc. (Bellcore)

Permission to use, copy, modify, and distribute this material for any purpose and without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies, and that the name of Bellcore not be used in advertising or publicity pertaining to this material without the specific, prior written permission of an authorized representative of Bellcore. BELLCORE MAKES NO REPRESENTATIONS ABOUT THE ACCURACY OR SUITABILITY OF THIS MATERIAL FOR ANY PURPOSE. IT IS PROVIDED "AS IS", WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES.

RSA Data Security, Inc., MD5 message-digest algorithm

Copyright© 1991-2, RSA Data Security, Inc. Created 1991. All rights reserved. License to copy and use this software is granted provided that it is identified as the "RSA Data Security, Inc. MD5 Message-Digest Algorithm" in all material mentioning or referencing this software or this function. License is also granted to make and use derivative works provided that such works are identified as "derived from the RSA Data Security, Inc. MD5 Message-Digest Algorithm" in all material mentioning or referencing the derived work. RSA Data Security, Inc. makes no representations concerning either the merchantability of this software or the suitability of this software for any particular purpose. It is provided "as is" without express or implied warranty of any kind. These notices must be retained in any copies of any part of this documentation and/or software.

mpack/munpack

Copyright© 1993,1994 by Carnegie Mellon University All Rights Reserved.

Permission to use, copy, modify, distribute, and sell this software and its documentation for any purpose is hereby granted without fee, provided that the above

copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of Carnegie Mellon University not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. Carnegie Mellon University makes no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

CARNEGIE MELLON UNIVERSITY DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS, IN NO EVENT SHALL CARNEGIE MELLON UNIVERSITY BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

TIFF

Copyright© 1988-1997 Sam Leffler Copyright© 1991-1997 Silicon Graphics, Inc.

Permission to use, copy, modify, distribute, and sell this software and its documentation for any purpose is hereby granted without fee, provided that (i) the above copyright notices and this permission notice appear in all copies of the software and related documentation, and (ii) the names of Sam Leffler and Silicon Graphics may not be used in any advertising or publicity relating to the software without the specific, prior written permission of Sam Leffler and Silicon Graphics.

THE SOFTWARE IS PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EXPRESS, IMPLIED OR OTHERWISE, INCLUDING WITHOUT LIMITATION, ANY WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. IN NO EVENT SHALL SAM LEFFLER OR SILICON GRAPHICS BE LIABLE FOR ANY SPECIAL, INCIDENTAL, INDIRECT OR CONSEQUENTIAL DAMAGES OF ANY KIND, OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER OR NOT ADVISED OF THE POSSIBILITY OF DAMAGE, AND ON ANY THEORY OF LIABILITY, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

PostgreSQL

Portions© 1996-2019, PostgreSQL Global Development Group Portions© 1994, The Regents of the University of California

Permission to use, copy, modify, and distribute this software and its documentation for any purpose, without fee, and without a written agreement is hereby granted, provided that the above copyright notice and this paragraph and the following two paragraphs appear in all copies.

IN NO EVENT SHALL THE UNIVERSITY OF CALIFORNIA BE LIABLE TO ANY PARTY FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, INCLUDING LOST PROFITS, ARISING OUT OF THE USE OF THIS SOFTWARE AND ITS DOCUMENTATION, EVEN IF THE UNIVERSITY OF CALIFORNIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

THE UNIVERSITY OF CALIFORNIA SPECIFICALLY DISCLAIMS ANY WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE SOFTWARE PROVIDED HEREUNDER IS ON AN "AS IS" BASIS, AND THE UNIVERSITY OF CALIFORNIA HAS NO OBLIGATIONS TO PROVIDE MAINTENANCE, SUPPORT, UPDATES, ENHANCEMENTS, OR MODIFICATIONS.

OpenSSL License

Copyright© 1998-2019 The OpenSSL Project. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment: "This product includes software developed by the OpenSSL Project for use in .the OpenSSL Toolkit." (<https://www.openssl.org/>)
4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact openssl-core@openssl.org.
5. Products derived from this software may not be called "OpenSSL" nor may "OpenSSL" appear in their names without prior written permission of the OpenSSL Project.
6. Redistributions of any form whatsoever must retain the following acknowledgment: "This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit." (<https://www.openssl.org/>)

THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT "AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product includes cryptographic software written by Eric Young (eyay@cryptsoft.com). This product includes software written by Tim Hudson (tjh@cryptsoft.com).

Original SSLeay License

Copyright© 1995-1998 Eric Young (eyay@cryptsoft.com) All rights reserved.

This package is an SSL implementation written by Eric Young (eyay@cryptsoft.com). The implementation was written so as to conform with Netscapes SSL.

This library is free for commercial and non-commercial use as long as the following conditions are adhered to. The following conditions apply to all code found in this distribution, be it the RC4, RSA, lhash, DES, etc., code; not just the SSL code. The SSL documentation included with this distribution is covered by the same copyright terms except that the holder is Tim Hudson (tjh@cryptsoft.com).

Copyright remains Eric Young's, and as such any Copyright notices in the code are not to be removed. If this package is used in a product, Eric Young should be given attribution as the author of the parts of the library used. This can be in the form of a textual message at program startup or in documentation (online or textual) provided with the package.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement: "This product includes cryptographic software written by Eric Young (eyay@cryptsoft.com)" The word "cryptographic" can be left out if the routines from the library being used are not cryptographic related :-).
4. If you include any Windows specific code (or a derivative thereof) from the apps directory (application code) you must include an acknowledgement: "This product includes software written by Tim Hudson (tjh@cryptsoft.com)"

THIS SOFTWARE IS PROVIDED BY ERIC YOUNG "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The licence and distribution terms for any publically available version or derivative of this code cannot be changed. i.e. this code cannot simply be copied and put under another distribution licence [including the GNU Public Licence.]

gawk

GNU GENERAL PUBLIC LICENSE Version 2, June 1991

<https://www.gnu.org/licenses/gpl-2.0.html>

Copyright© 1989, 1991 Free Software Foundation, Inc.

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

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

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Lesser General Public License instead.) You can apply it to your programs, too.

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

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

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

1. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or

translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

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

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

3. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:
 - a. You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.
 - b. You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.
 - c. If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.) These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

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

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

4. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:
 - a. Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
 - b. Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
 - c. Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.) The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

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

5. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.
6. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.
7. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.
8. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

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

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

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

9. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.
10. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

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

11. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally. NO WARRANTY
12. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.
13. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Ghostscript

The files in the base, psi, lib, toolbin, examples, doc and man directories (folders) and any subdirectories (sub-folders) thereof are part of GPL Ghostscript.

The files in the Resource directory and any subdirectories thereof are also part of GPL Ghostscript, with the explicit exception of the files in the CMap subdirectory (except "Identity-UTF16-H", which is part of GPL Ghostscript). The CMap files are copyright Adobe Systems Incorporated and covered by a separate, GPL compatible license.

The files under the jpegxr directory and any subdirectories thereof are distributed under a no cost, open source license granted by the ITU/ISO/IEC but it is not GPL compatible - see jpegxr/COPYRIGHT.txt for details.

GPL Ghostscript is free software; you can redistribute it and/or modify it under the terms the GNU General Public License as published by the Free Software Foundation, either version 3 of the License, or (at your option) any later version.

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

You should have received a copy of the GNU General Public License along with this program so you can know your rights and responsibilities. It should be in a file named doc/COPYING. If not, write to the

Free Software Foundation, Inc.,
59 Temple Place Suite 330,
Boston, MA, 02111-1307, USA.

GPL Ghostscript contains an implementation of techniques covered by US Patents 5,055,942 and 5,917,614, and corresponding international patents. These patents are licensed for use with GPL Ghostscript under the following grant:

Whereas, Raph Levien (hereinafter "Inventor") has obtained patent protection for related technology (hereinafter "Patented Technology"), Inventor wishes to aid the the GNU free software project in achieving its goals, and Inventor also wishes to increase public awareness of Patented Technology, Inventor hereby grants a fully paid up, nonexclusive, royalty free license to practice the patents listed below ("the Patents") if and only if practiced in conjunction with software distributed under the terms of any version of the GNU General Public License as published by the

Free Software Foundation, 59 Temple Place, Suite 330, Boston, MA 02111.

Inventor reserves all other rights, including without limitation, licensing for software not distributed under the GNU General Public License.

5055942 Photographic image reproduction device using digital halftoning to para images allowing adjustable coarseness 5917614 Method and apparatus for error diffusion paraing of images with improved smoothness in highlight and shadow regions

MariaDB and FreeTDS

Copyright© 1991, 1999

Free Software Foundation, Inc.
51 Franklin Street, Fifth Floor, Boston, MA 02110-1301, USA

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

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

Preamble The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public Licenses are intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users.

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

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

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

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

We protect your rights with a two-step method:

1. we copyright the library, and
2. we offer you this license, which gives you legal permission to copy, distribute and/or modify the library.

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

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

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

When a program is linked with a library, whether statically or using a shared library, the combination of the two is legally speaking a combined work, a derivative of the original library. The ordinary General Public License therefore permits such linking only if the entire combination fits its criteria of freedom. The Lesser General Public License permits more lax criteria for linking other code with the library.

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

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

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

Although the Lesser General Public License is Less protective of the users' freedom, it does ensure that the user of a program that is linked with the Library has the freedom and the wherewithal to run that program using a modified version of the Library.

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

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

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

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

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

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

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

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

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

3. You may modify your copy or copies of the Library or any portion of it, thus forming a work based on the Library, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:
 - a. The modified work must itself be a software library.
 - b. You must cause the files modified to carry prominent notices stating that you changed the files and the date of any change.
 - c. You must cause the whole of the work to be licensed at no charge to all third parties under the terms of this License.
 - d. If a facility in the modified Library refers to a function or a table of data to be supplied by an application program that uses the facility, other than as an argument passed when the facility is invoked, then you must make a good faith effort to ensure that, in the event an application does not supply such function or table, the facility still operates, and performs whatever part of its purpose remains meaningful. (For example, a function in a library to compute square roots has a purpose that is entirely well-defined independent of the application. Therefore, Subsection 2d requires that any application-supplied function or table used by this function must be optional: if the application does not supply it, the square root function must still compute square roots.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Library, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Library, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

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

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

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

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

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

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

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

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

However, linking a "work that uses the Library" with the Library creates an executable that is a derivative of the Library (because it contains portions of the Library), rather than a "work that uses the library". The executable is therefore covered by this License. Section 6 states terms for distribution of such executables.

When a "work that uses the Library" uses material from a header file that is part of the Library, the object code for the work may be a derivative work of the Library even though the source code is not. Whether this is true is especially significant if the work can be linked without the Library, or if the work is itself a library. The threshold for this to be true is not precisely defined by law.

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

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

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

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

Also, you must do one of these things:

- a. Accompany the work with the complete corresponding machine-readable source code for the Library including whatever changes were used in the work (which must be distributed under Sections 1 and 2 above); and, if the work is an executable linked with the Library, with the complete machine-readable "work that uses the Library", as object code and/or source code, so that the user can modify the Library and then relink to produce a modified executable containing the modified Library. (It is understood that the user who changes the contents of definitions files in the Library will not necessarily be able to recompile the application to use the modified definitions.)
- b. Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that
 - i. uses at run time a copy of the library already present on the user's computer system, rather than copying library functions into the executable, and
 - ii. will operate properly with a modified version of the library, if the user installs one, as long as the modified version is interface-compatible with the version that the work was made with.
- c. Accompany the work with a written offer, valid for at least three years, to give the same user the materials specified in Subsection 6a, above, for a charge no more than the cost of performing this distribution.
- d. If distribution of the work is made by offering access to copy from a designated place, offer equivalent access to copy the above specified materials from the same place.
- e. Verify that the user has already received a copy of these materials or that you have already sent this user a copy. For an executable, the required form of the "work that uses the Library" must include any data and utility programs needed for reproducing the executable from it. However, as a special exception, the materials to be distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

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

- f. You may place library facilities that are a work based on the Library side-by-side in a single library together with other library facilities not covered by this License, and distribute such a combined library, provided that the separate distribution of the work based on the Library and of the other library facilities is otherwise permitted, and provided that you do these two things:
 - a. Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities. This must be distributed under the terms of the Sections above.
 - b. Give prominent notice with the combined library of the fact that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.
- g. You may not copy, modify, sublicense, link with, or distribute the Library except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, link with, or distribute the Library is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.
- h. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Library or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Library (or any work based on the Library), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Library or works based on it.
- i. Each time you redistribute the Library (or any work based on the Library), the recipient automatically receives a license from the original licensor to copy, distribute, link with or modify the Library subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties with this License.
- j. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Library at all. For example, if a patent license would not permit royalty-free redistribution of the Library by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Library.

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

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

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

- k. If the distribution and/or use of the Library is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Library under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.
- l. The Free Software Foundation may publish revised and/or new versions of the Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Library specifies a version number of this License which applies to it and "any later

version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Library does not specify a license version number, you may choose any version ever published by the Free Software Foundation.

- m. If you wish to incorporate parts of the Library into other free programs whose distribution conditions are incompatible with these, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

- n. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE LIBRARY IS WITH YOU. SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.
- o. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

QtAV

© Wang Bin wbsecg1@gmail.com Shanghai University->S3 Graphics->Deepin, Shanghai, China 2013-01-21

QtAV is free software licensed under the term of LGPL v2.1. The player example is licensed under GPL v3. If you use QtAV or its constituent libraries, you must adhere to the terms of the license in question.

Rather than repeating the text of the LGPL v2.1, the original text can be found in [GNU LESSER GENERAL PUBLIC LICENSE, Version 2.1](#).

FFmpeg

Most files in FFmpeg are under the GNU Lesser General Public License version 2.1 or later (LGPL v2.1+). Read the file `COPYING.LGPLv2.1` for details. Some other files have MIT/X11/BSD-style licenses. In combination the LGPL v2.1+ applies to FFmpeg.

Rather than repeating the text of the LGPL v2.1, the original text can be found in [GNU LESSER GENERAL PUBLIC LICENSE, Version 2.1](#).

c3.js

The MIT License (MIT) © 2013 Masayuki Tanaka

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

d3.js

Copyright© 2010-2017 Mike Bostock All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. * Neither the name of the author nor the names of contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

jwt-cpp

MIT License

Copyright © 2018 Dominik Thalhammer

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

QXlsx

The MIT License

Copyright © 2017-, <https://github.com/j2doll/QXlsx>

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.