

Introduction

Welcome to Project File Systems! In this project, you will add features to the file system of Pintos.

In Project User Programs, you implemented much of the syscall functionality for the file system syscalls. However, much of the internals were abstracted away from you as you called existing file system functions. In this project, you will dive deep into the file system all the way down to the bytes that are stored on disk. In addition, you will also work on adding a buffer cache to speed up accesses to disk.

The details of this assignment can be found in the [Tasks](#) section. However, you may find it helpful to first read through some of the [CS 162 Pintos documentation](#). This will likely help you understand the required tasks.

Even if you read through some of these during previous projects, we recommend (re)reading through the following sections of the [Pintos documentation](#):

- [File systems](#)

Setup

For this project, you will be building off of your Project User Programs solutions.

First, log into your VM. Similar to what you did when starting Project Threads, we recommend that you tag your final Project Threads code in case you ever want to refer back to it.

```
cd ~/code/group
git tag proj-threads-completed
git push group main --tags
```

Next, restore your code to your Project User Programs solution.

```
git fetch --all --tags --jobs=2
rm -rf src/
git pull group main
```

```
git checkout proj-userprog-completed -- src/  
git commit -m "Restore project userprog solution"  
git push group main  
git pull staff main
```