

# Project Description: Personal Virtual File System

ETH Zurich

April 20, 2013

Hand-out: April 20, 2013  
Due: May 13, 2013 (final deadline)

## 1 Overview

The goal of this project is to develop a single-file Virtual File System (**VFS**) in either Java or C#. A VFS is usually built on top of a host file system, and it enables access to files located in different host file systems in a uniformed way. A single-file VFS could also be used as a manageable container with the functionality of a concrete file system through the usage of software. Typical applications of a single-file VFS include virtual machines, emulators, encryption (providing a FS with encrypted data), or even programming IDEs (IntelliJ<sup>1</sup>). You can find more information on VFSs on Wikipedia<sup>2</sup>.

## 2 Deadlines

All project work must be done in teams of **3 people**. The project itself is divided in 3 parts, each with its own deadline:

- The VFS core (due April 8, 2013)
- The VFS browser (due April 22, 2013)
- The VFS synchronization server (due May 13, 2013)

At each deadline you need to hand in a written report on the design and implementation of the corresponding part. The proposed project structure allows incremental development.

## 3 Code Repository

Use the SVN repository at <https://code.vis.ethz.ch/>. Create a project using as name the groupID that will be sent to you via email. Add your assistants username as project member (for usernames see 7).

## 4 Project Grading

You can earn 100 points in the project. The points are composed as follows:

- Requirements coverage, max 40 points
- Code quality, max 20 points

---

<sup>1</sup><http://confluence.jetbrains.com/display/IDEADEV/IntelliJ+IDEA+Virtual+File+System>

<sup>2</sup>[http://en.wikipedia.org/wiki/Virtual\\_file\\_system](http://en.wikipedia.org/wiki/Virtual_file_system)

- Unit tests, max 15 points
- Final report, max 15 points
- Bonus features, max 10 points<sup>3</sup>

The grade will be determined based on the project's final implementation and the final report as submitted by May 13. However, each team has to provide an implementation and a report summarizing the current design and project status at each intermediate deadline. The implementation at the intermediate deadlines has to be runnable.

## 4.1 Statement coverage of unit tests

Statement coverage<sup>4</sup> is a measure used in software testing to describe the percentage of statements that has been tested: in general, the higher the coverage, the more reliable the test results. In this project we will use the following tools to collect coverage information on your tests.<sup>5</sup>

- Java:  
Eclipse and the EclEmma plugin (<http://www.eclEmma.org/>)
- C#: JetBrains dotCover (<http://www.jetbrains.com/dotcover/>)<sup>6</sup>

## 4.2 Code smells

Code smells<sup>7</sup> are indications of potential problems in the source code. We will use the following tools to detect smells in your code:

- Java: PMD (<http://pmd.sourceforge.net/>)
- C#: VS Code Analysis (<http://msdn.microsoft.com/en-us/library/3z0aeatx.aspx>)

For both tools the quick start guide is available on the course's website.

## 5 General notes

1. The VFS core, except for the bonus features, may *not* rely on any DBMS system or 3rd party library.
2. The system should be robust with respect to incorrect user input. For example, when a user tries to import files from a nonexistent directory, the VFS should not crash.
3. Design your application in a modular way to support separation of concerns. For example, VFS core should not depend on the GUI.

## 6 Project Details

### 6.1 Part 1: VFS core (due April 8)

The goal of the first part is to develop the **VFS core**. The VFS core operates on **virtual disks**, each being a single file in the host file system. The core provides an API to facilitate creating, modifying and deleting virtual disks. Particularly, the core provides means for its client to: 1) create a virtual

---

<sup>3</sup>There's a total of 46 bonus points available. The points listed on each bonus feature are the *maximum* number of points that can be achieved for that feature. If you achieve more than 10 bonus points in total, the remaining points can be used to compensate missing points in other areas of the projects (but not the exam).

<sup>4</sup>[http://en.wikipedia.org/wiki/Code\\_coverage](http://en.wikipedia.org/wiki/Code_coverage)

<sup>5</sup>All tools will be running under the default settings.

<sup>6</sup>You can use dotCover with the academic license we sent you before.

<sup>7</sup>[http://en.wikipedia.org/wiki/Code\\_smell](http://en.wikipedia.org/wiki/Code_smell)

disk of user-specified size, as well as delete an existing virtual disk; 2) import files and directories from the host file system to a virtual disk, as well as export files and directories from the VFS to the host file system; and 3) navigate through the directories of a virtual disk, as well as rename, copy, remove or move existing files and directories in the virtual disk. The size of a virtual disk is fixed during its lifetime, and errors should be reported when pending operations are infeasible due to this restriction.

### 6.1.1 Requirements

1. The virtual disk must be stored in a single file in the working directory in the host file system.
2. VFS must support the creation of a new disk with the specified maximum size at the specified location in the host file system.
3. VFS must support several virtual disks in the host file system.
4. VFS must support disposing of the virtual disk.
5. VFS must support creating/deleting/renaming directories and files.
6. VFS must support navigation: listing of files and folders, and going to a location expressed by a concrete path.
7. VFS must support moving/copying directories and files, including hierarchy.
8. VFS must support importing files and directories from the host file system.
9. VFS must support exporting files and directories to the host file system.
10. VFS must support querying of free/occupied space in the virtual disk.

### 6.1.2 Bonus Features

#### Basic

1. Compression, if implemented with 3d party library. (1p)
2. Encryption, if implemented with 3d party library. (1p)
3. Elastic disk: Virtual disk can dynamically grow or shrink, depending on its occupied space. (2p)

#### Advanced

1. Compression, if implemented by hand (you can take a look at the arithmetic compression<sup>8</sup>). (2p)
2. Encryption, if implemented by hand. (3p)
3. Large data: This means, that VFS core can store & operate amount of data, that can't fit to PC RAM ( typically, more than 4Gb). (3p)

### 6.1.3 Hints

1. Develop a set of unit tests prior to the development, try to work using the Test-Driven-Development (TDD) approach.
2. The key challenge is to develop an efficient storage structure at physical layer. You can take an advantage of specialized data structures, like B-trees<sup>9</sup>.
3. Keep in mind that in the next part you will be required to implement file search.
4. At this point, it is helpful to develop a bash-like console application, to test your VFS (with operations like ls, cd, copy, etc)

---

<sup>8</sup>[http://en.wikipedia.org/wiki/Arithmetic\\_coding](http://en.wikipedia.org/wiki/Arithmetic_coding)

<sup>9</sup>[http://en.wikipedia.org/wiki/B-tree#In\\_filesystems](http://en.wikipedia.org/wiki/B-tree#In_filesystems)

## 6.2 Part 2: VFS browser (due April 22)

The goal of this part is to provide a Graphical User Interface (GUI) for the VFS core. The target platform is up to you: it could be a desktop application (Windows or Linux), web-application or a mobile client. The GUI must support all implemented operations from Part 1. We require navigation with both mouse/touch and keyboard. For this part of the project you only have to support a single user of the VFS, multi-user support is not required. You are encouraged to design a nice and user-friendly interface. Bonus points are given for a responsive GUI that does not stall during long-running operations, like import or search. Programming an additional, different client is another option to earn bonus points.

### 6.2.1 Requirements

1. The browser should be implemented on one of the following platforms: desktop, web or mobile.
  - (a) Web applications written in C# must use Silverlight.
  - (b) Web applications written in Java must either use Google Web Toolkit or JavaFX.
  - (c) Mobile applications must target Android or Windows Phone.
2. The browser should support all operations from Part 1 (VFS core). For example, users should be able to select a file/folder and copy it to another location without using console commands.
3. The browser should support both single and multiple selection of files/folders.
4. The browser should support keyboard navigation. The mandatory set of operations includes folder navigation, going to parent and child folders (this is optional for mobile applications due to limited keyboard functionality).
5. The browser should support mouse navigation (or touch in case of the mobile platform). The required operations are the same as in requirement 4.
6. The browser should support file-name search based on user-given keywords. The search should provide options for: case sensitive/ case insensitive search; restrict search to folder; restrict search to folder and subfolders.

### 6.2.2 Bonus Points

#### Basic

1. Responsive UI, i.e. the browser does not stall during long-running operations (i.e. file search or import). (3p)
2. Advanced search; For example, search with wildcards/regexp, and approximate search based on some metric, e.g. *edit distance*. (2p)
3. Nice-to-have features like operation progress report (e.g. the number of files processed during export) or drag-and-drop for manipulative operations (move, copy, import). (2p)

#### Advanced

1. The browser is implemented for an additional platform. (5p)
2. Efficient full-text search (using some sort of indexing). (4p)

Other additional features can be approved by your assistant on a case-by-case basis.

### 6.2.3 Hints

1. The web interface of Dropbox could be used as an example of a good interface.

### 6.3 Part 3: Synchronization Server (due May 13)

The goal of this part is to develop a synchronization server which will propagate changes in a virtual disk from one machine to another using a network connection. The management of the distributed VFSs is made on an account basis, i.e. in order to use the server, a user must have an account and link a virtual disk to this account. The synchronization should support a single using a virtual disk on multiple machines. It does not need to support synchronization between multiple users of the same disk. Note that unlinked disks should not be synchronized, and that you do not have to implement a conflict resolution strategy (e.g. for the case where a document is changed on two different machines at the same time). You should extend the browser from the previous part to be a client of the synchronization server. Additional points are given for implementing a concurrent server (when changes on the same account are done simultaneously on different machines).

#### 6.3.1 Requirements for the browser

1. The browser should allow the user to create a new account or to log in to an existing account.
2. The browser should offer to switch to an offline mode, and be able to operate without a connection to the server.
3. The browser should support binding an existing virtual disk to an active account.

#### 6.3.2 Requirements for the server

1. The server should support registration of unique accounts. Each account includes name & password.
2. The server should track changes to linked virtual disks of registered accounts and synchronize the changes across the machines.

#### 6.3.3 Bonus Points

##### Basic

1. Provide a set of mocked unit tests for your implementation. (1p)
2. Conflict Resolution: implement a conflict resolution scheme, so that concurrent changes to the same file are not lost (e.g. saving conflicting files as separate versions). (3p)

##### Advanced

1. A concurrent client/server system:
  - The browser is updating automatically when changes to the disk occur. (3p)
  - The server is able to synchronize changes, which are done simultaneously on the same account on different machines. (3p)
2. Incremental Changes: minimize the communication between the browser and the server by only transferring the parts of a file that changed. (4p)
3. File History: provide a history for files and an interface to restore a previous version. (4p)

## 7 Contacts

Christian Estler (estlerh)	christian.estler@inf.ethz.ch
Alexey Kolesnichenko (alexeyk)	alexey.kolesnichenko@inf.ethz.ch
Max (Yu) Pei (ypei)	yu.pei@inf.ethz.ch
Julian Tschannen (juliant)	julian.tschannen@inf.ethz.ch