

Distributed Storage

Evan Danaher

1 Introduction

In many computer labs, there are large numbers of computers with unused drive space. There may also be relatively large quantities of data to be backed up. The goal of this project was to develop a system for storing data distributed over many computers, with enough redundancy so that data can still be recovered if several of the machines are unavailable (due to inevitable hardware failure). The RSraid scheme was chosen for maximum flexibility, and the result is an extremely robust, though very user unfriendly backup system.

2 Description

The problem of effectively backing up critical data has been around for many decades. Traditional backup methods use a tape drive or other device to store data separate from the main computer. However, this requires the purchase of an extra drive and extra media to use for backups, as well as human labor to ensure that tapes are stored safely and not overwritten. A different method has become possible in the past few years, due to the prevalence of large hard drives with significant unused space in standard desktops and networking of these computers. This unused space can be used for automatically backing up data at no additional expense. However, since desktops are often less reliable, this backup method requires data to be recoverable, even if a significant amount is lost. Redundancy techniques provide this ability.

An article, *A Tutorial on Reed-Solomon Coding for Fault-Tolerance in RAID-like Systems*, describes the use of a variant of Raid-Solomon for use in certain applications. The idea is that Raid-Solomon can be used relatively simply when you know which data is lost; normally it is used in devices such as CD-ROMS, where it is impossible to know which data is correct and which was mangled. But in "RAID-like Systems," it is known which device failed, so data recovery is simpler. In particular, I can use this algorithm for recovering data from multiple computers when some computers are dead.

This project is designed for a fairly limited but still significant group: organizations that have large file servers to store most data, as well as dozens or hundreds of "user" machines that are relatively new. Only these groups will be able to use this program: with no central server to back up, there is no need for it; without a number of "user" machines, there is not enough space to store the backup. However, this category still contains many groups, such as businesses and schools. Even large businesses with existing backup systems may find this useful; once set up, it requires little maintenance, and can recover backups quickly compared with existing solutions.