

Performance Comparison of File Systems on a Virtual Guest

Jacob Carney and Dr. Robert Marmorstein (Longwood University)

Introduction

A **virtual machine** is a program that simulates the hardware of a computer system in software. Virtual machines are the key to **cloud computing**, since a single powerful server can run multiple virtual machines without allowing programs running on one to affect the others. The computer a virtual machine runs on is called a **host machine**. The operating system that runs inside the virtual machine is called the **guest machine** or **virtual guest**.

One key factor in the efficiency of a virtual machine is the choice of a file system. A **file system** is a data structure that manages the organization of small blocks of data into files on the disk..

Modern operating systems support several different file systems. Some of the most popular file systems are **ext4** (the default file system in Linux), **ntfs** (used by Windows computers), **btrfs** (an alternative to ext4), and **xfs** (a system designed by SGI for high performance servers). Choice of a file system can have a significant impact on the speed of a server..

There are two file systems used by a virtual machine: the file system on the host machine (the **host file system**) and the file system on the guest machine (the **guest file system**). The host file system is stored directly on the hard drive of the server. The guest file system is stored in a special file within the host file system called a **disk image**.

While the choice of host file system has been well studied, the performance effects of using different guest file systems has not. In this project, we compare the performance of virtual machines using different guest file systems.

Technical Background

One of the most common programs for hosting virtual machines is VirtualBox, a program released by Oracle. Virtualbox stores files using the VDI disk image format. A VDI file is divided into four parts: a header, an image block map, block alignment padding, and image blocks which store the actual data of the files [1].

Implementation

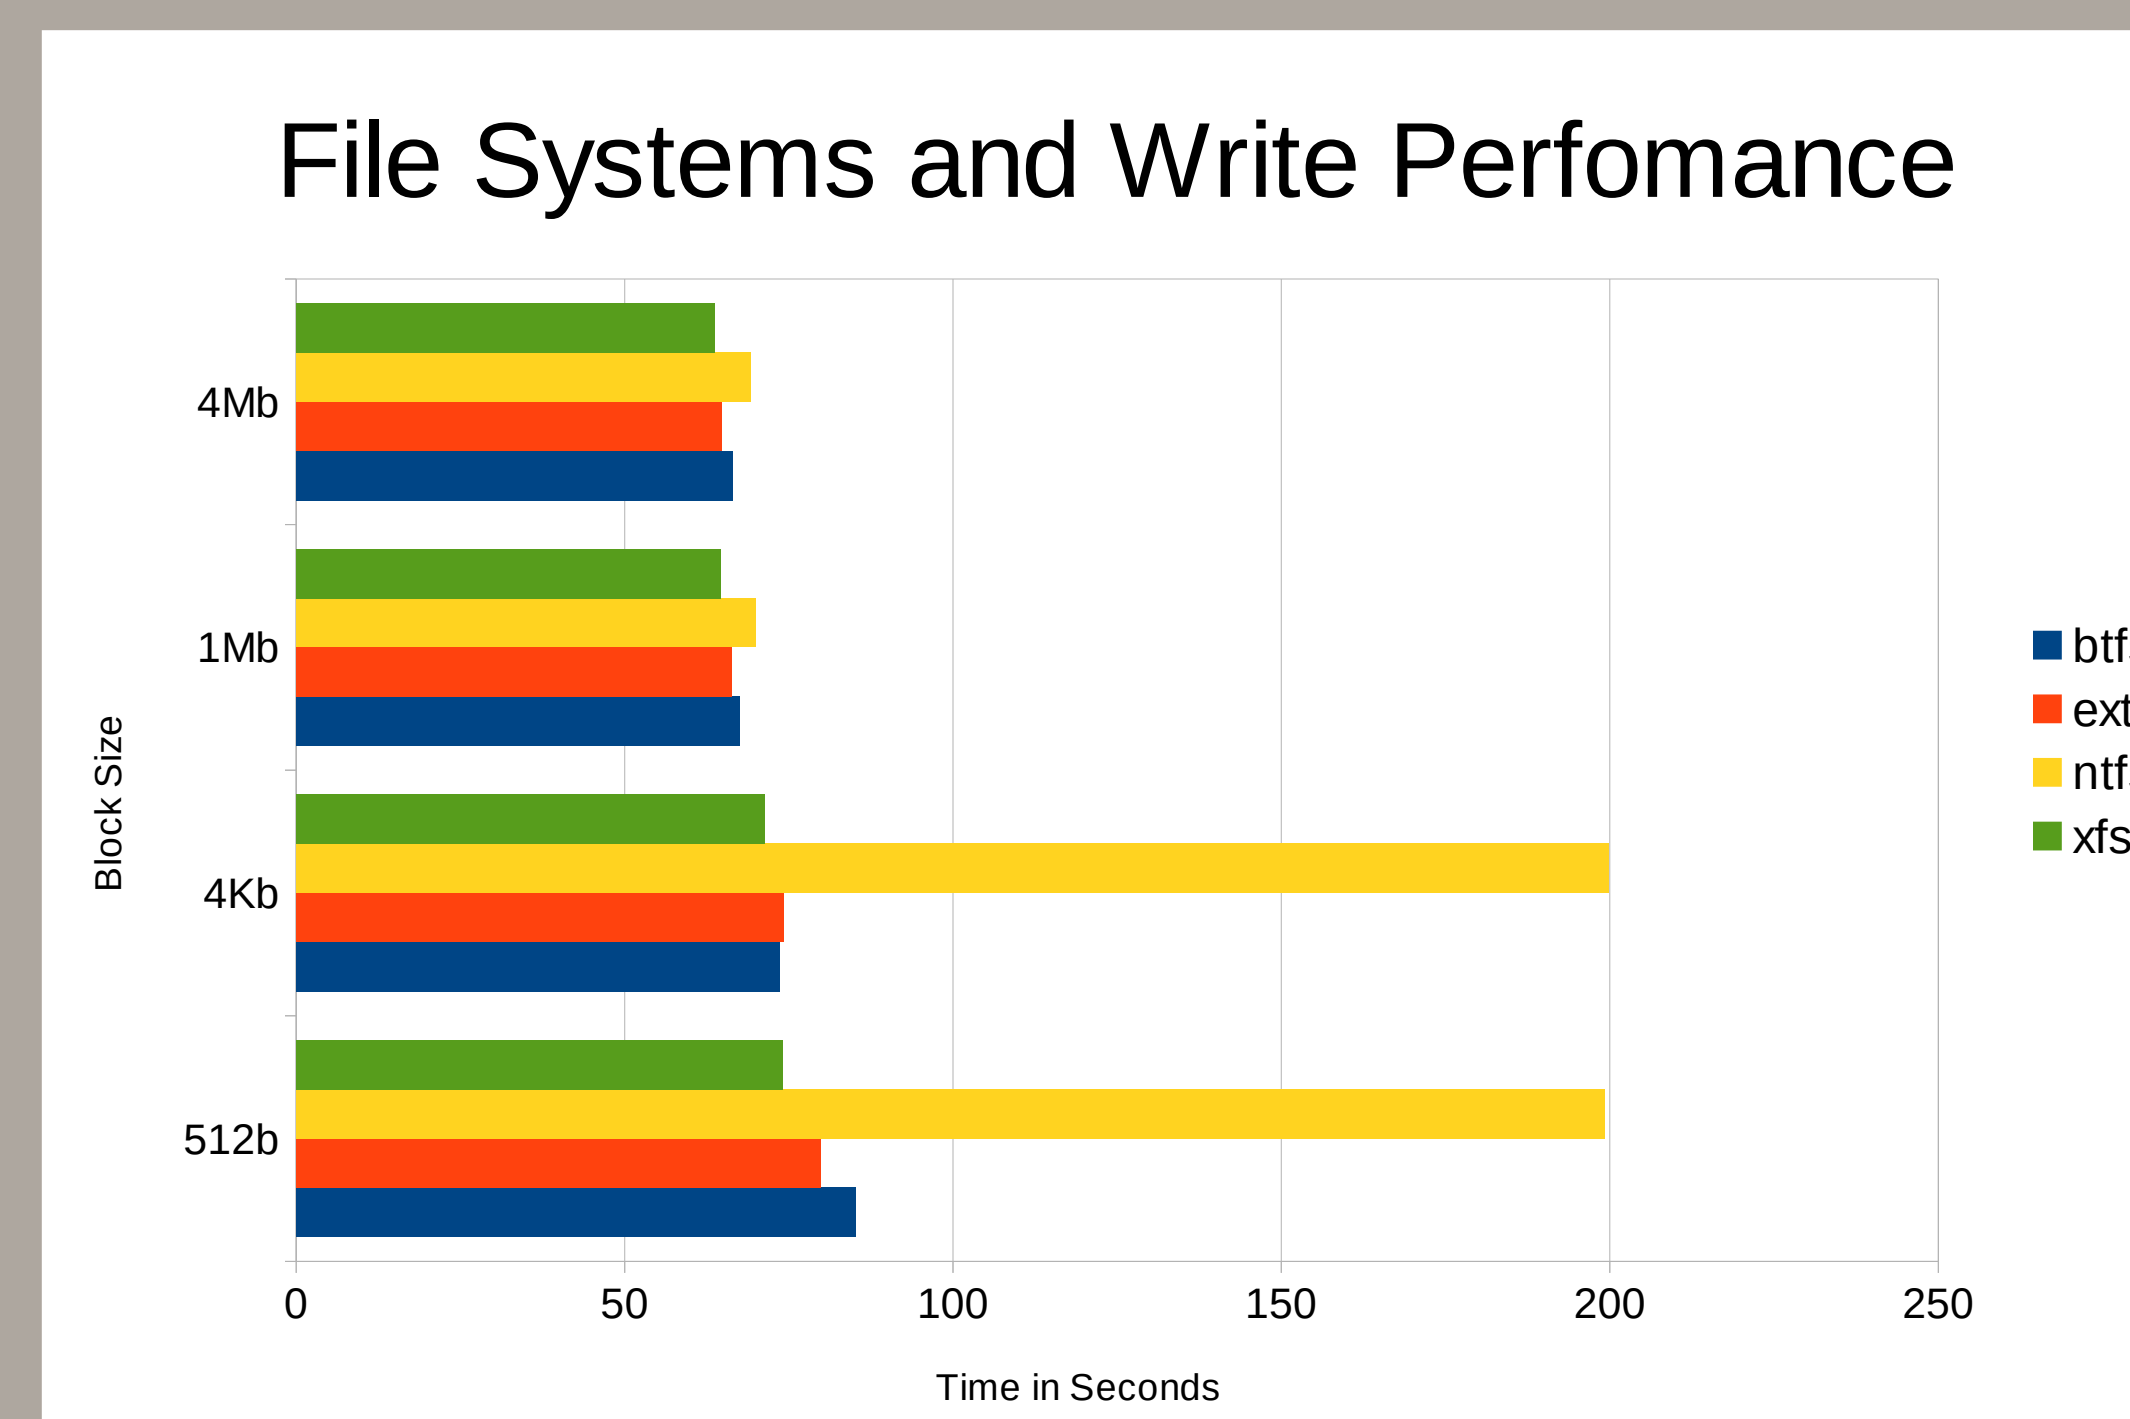
To compare performance of different guest file systems, we created four CentOS Linux virtual machine images containing a 20Gb data partition. We formatted the data partition of each machine with a different file system. Then we used a tool called Vagrant to distribute all four file systems to twenty-four Linux workstations.

Using Unix shell scripts (shown below), we logged in to each system and launched a program which wrote a large (1Gb) file to the data partition. By timing the length of time required to create the file, we measured the write speed on the virtual system under each file system. We performed 30 trials of each test and recorded the times using the Unix “time” command.

Because block size can have a significant impact on write performance, we repeated this test using several different write sizes. We used strides of 512 bytes, 4Kib, 1Mib, and 4Mib.

To mitigate variance caused by caching, we performed a priming run to “warm up the caches” before collecting data.

Data



Conclusion

The data shows that XFS has a slight advantage over the other file systems across all block sizes and ntfs has performance issues on smaller block sizes.

There are many ways this work could be extended. This project focused on write performance, but for many workloads, read performance is critical. We could also examine performance creating larger or smaller data sets, using different strides.

The file created by this project consisted of a series of random numbers. Examining performance on a more realistic workload, such as compilation of the Linux kernel or the gcc compiler or simulation of a realistic web server might be of interest.

We also hope to consider how choice of the host file system affects performance of the guest file system.

In addition, while ext4, btrfs, ntfs, and xfs are very commonly used filesystems, there are many others that might be worth examining, such as JFS, ReiserFS, ZFS, VFAT, and older versions of the Linux ext filesystem.

Reference

[1] “All About VDIs”, Terry E, virtualbox.org, website, 28 May 2008.