Study of file system performance varying the operating environment

Ruben Pinilla

Computer Architecture Department, Technical University of Catalonia
c/ Jordi Girona, 1-3, Campus Nord, 08034 Barcelona, Spain

e-mail: rpinilla@ac.upc.es

Abstract

This work studies the file system performance variation when the operating environment is changed. We have developed some experiments based on the IOzone benchmark to characterize the overhead introduced when using a host operating system running over a VMware virtual machine.

Experiments results indicate that the overhead introduced by the VMM is generalized on the I/O operations. Depending on the operating environment, we have found that some write operations run faster over a VMware VM than running on the base OS.

Keywords: Operating Environments, VMware, file system performance, virtual machine.

1. INTRODUCTION

In this work we present an analysis of the overhead introduced by using a VMM (Virtual Machine Monitor) when managing a file system. The added VMM abstracts underlying operating system functionalities and offers a hardware platform emulation to the host operating system.

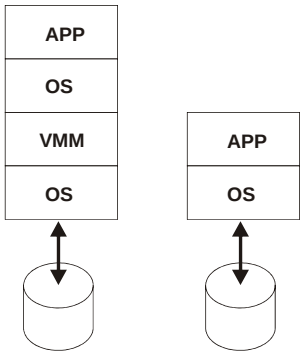

Figure 1. Operating environments to test.

* This work has been supported by the Ministry of Science and Technology of Spain and by the European Union (FEDER) under contract TIC2001-0995-C02-01. All products or company names mentioned herein are the trademarks or registered trademarks of their respective owners.

The main goal of this work is to determine the overhead that suffers a VMM hosted operating system managing the file system compared to the same operating system running directly over the hardware platform.

The rest of the document is structured as follows: section 2 presents some experimental results to characterize the performance of file system operations varying the operating environment and section 3 concludes the paper.

2. EVALUATION

2.1. Experiments

We have developed a set of experiments using the same hardware platform based on a Intel 4 Xeon 2,8 GHz with 1 Gbytes of RAM. The I/O subsystem is composed by a LSI 320 Mb/s SCSI controller with two 36 Gbytes SCSI disks. As VMM we have used VMware Workstation release 4.5.0 [1] running on top of Windows XP Professional and Linux Red Hat 9.0 operating systems. In addition to the base operating systems configurations, the operating environments used are explained in Table 1.

Table 1. Operating environments used in experiments.

Operating Environment	Host OS	Base OS
OE1	Windows XP Professional SP1	Windows XP Professional SP1
OE2	Windows XP Professional SP1	Linux Red Hat 9.0 Kernel 2.4.20
OE3	Linux Red Hat 9.0 Kernel 2.4.20	Windows XP Professional SP1
OE4	Linux Red Hat 9.0 Kernel 2.4.20	Linux Red Hat 9.0 Kernel 2.4.20

The base operating system installation has been done using one disk for each OS using the default settings during their setup (the Red Hat 9.0 setup uses the default workstation settings). The host operating system installation has been realized creating two VMware virtual machines using the same base operating system settings. Both VMs have been created with the next settings:

- 256 Mbytes of RAM.
- Virtual disk type: SCSI.
- Virtual disk size: 5 Gbytes.

To determine the performance of the file system on each configuration, we have measured the transfer rate of a set of file operations based on the *IOzone file system benchmark* [2]:

Table 2. File operations considered in experiments.

File Operation	Description
Read	Measures the performance of reading an existing file
Re-read	Measures the performance of reading a file that was recently read
Write	Measures the performance of writing a new file
Re-write	Measures the performance of writing a file that already exists
Fread	Measures the performance of reading a file using the library function fread()
Freread	Measures the performance of reading a file using the library function fread() that was recently read
Fwrite	Measures the performance of writing a file using the library function fwrite()
Frewrite	Measures the performance of writing an existing file using the library function fwrite()

We have executed the *IOzone benchmark* to collect measures of I/O transfer rates as follows:

Table 3. Parameters used in the benchmark execution.

`iozone -Rab read-write.wks -g 128m -i 0 -i 1`
Runs the read, reread, write and re-write tests and outputs the results to file read-write.wks. The size of the file can take values from 64 Kbytes to 128 Mbytes. The record size varies from 4 Kbytes to 16384 Kbytes.
`iozone -Rab fread-fwrite.wks -g 128m -i 6 -i 7`
Runs the fread, freread, fwrite and frewrite tests and outputs the results to file fread-fwrite.wks. The size of the file can take values from 64 Kbytes to 128 Mbytes. The record size varies from 4 Kbytes to 16384 Kbytes.

The set of parameters that are involved in the experiments, depending on the file operation, are the following:

- Read transfer rate (Bytes read / second)
- Write transfer rate (Bytes written / second)

2.2. Results

We have executed the *IOzone benchmark* as described in Table 3 over the next configurations:

Experiment 1

This experiment contemplates two cases, depending on the OS to compare:

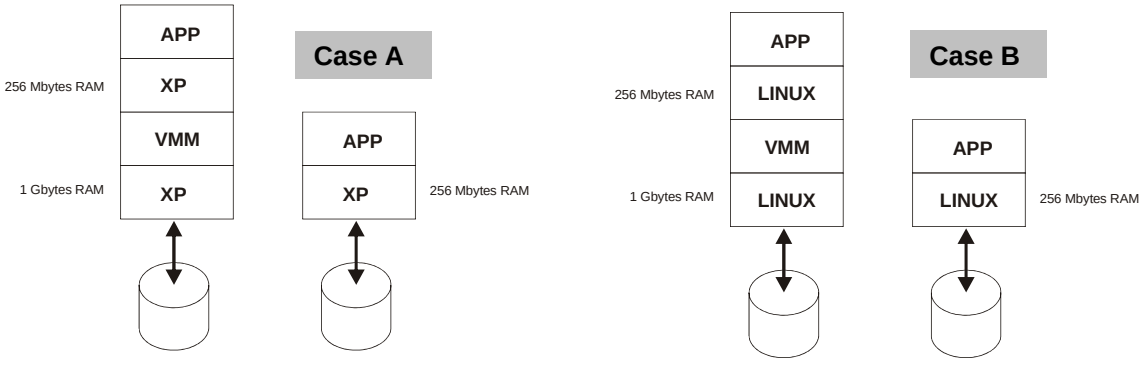

Figure 2. Operating environments to test in experiment 1. A) Compare Windows XP. B) Compare Linux.

The goal of this experiment is to detect possible performance differences when a VMM is added over an operating system. To quantify the overhead introduced by the use of VMware, we have run the operations benchmarks from Table 2 over the scenarios described in Figure 2.

In case A, we observe a better performance in the *write* operation (see Table 4) on the host operating system as well as in the *rewrite* operation but with less difference. On the other hand, *read* and *reread* operations perform better on the base OS. The results obtained using the library operations show a similar behavior.

In case B, we have found that Linux behaves better when it runs directly over the hardware. Only for file size of 128 Mbytes we detect superior performance for the VMware configuration. As occurs in case A, library operations show comparable behavior.

We include in Table 4 some representative values of the resulting speedup calculated dividing the base OS measures by the host OS measures on the *write* operation. Values less than 1 indicates that host OS performs better that the base OS and values greater that 1, the base OS performs better. We can observe that for case A the host OS configuration presents an enhanced performance and for case B is the base OS that gets the best results.

Table 4. Speedups calculated dividing the base OS measures by the host OS measures on the write operation (case A: Windows XP, case B: Linux).

Case A		Record size (Kbytes)				
		4	8	16	32	64
File size (Kbytes)	64	0,81	1,26	0,37	1,11	0,12
	128	1,01	0,18	0,17	1,58	1,00
	256	0,27	0,99	1,23	0,95	0,93
	512	0,90	1,06	0,91	0,94	1,31
	1024	2,24	1,02	0,93	0,92	0,91
	2048	0,70	0,93	0,72	0,93	0,94
	4096	0,93	0,92	0,94	0,89	0,94
	8192	0,95	0,92	0,96	0,95	0,94
	16384	0,93	0,92	1,01	0,96	0,93

Case B		Record size (Kbytes)				
		4	8	16	32	64
File size (Kbytes)	64	0,01	2,94	2,96	3,21	3,64
	128	3,82	2,24	2,31	1,96	1,99
	256	1,26	1,99	1,91	1,69	1,80
	512	3,11	7,05	1,70	1,60	1,57
	1024	4,06	1,72	1,56	4,06	4,18
	2048	2,38	2,89	3,20	0,47	1,15
	4096	2,91	0,80	0,91	0,78	1,42
	8192	2,54	1,91	2,43	1,93	1,06
	16384	2,56	1,58	1,34	1,24	1,13

In Figure 3, we show a graph representing the speedup on the *read* operation obtained dividing the base OS measures by the host OS measures. Values less than 1 indicates that host OS performs better that the base OS. On the contrary, values greater that 1, the base OS performs better. For case A, there is an overall uniform performance behavior, but in the case B there are situations where the host OS has superior results than the base Linux configuration (file size of 128 Mbytes).

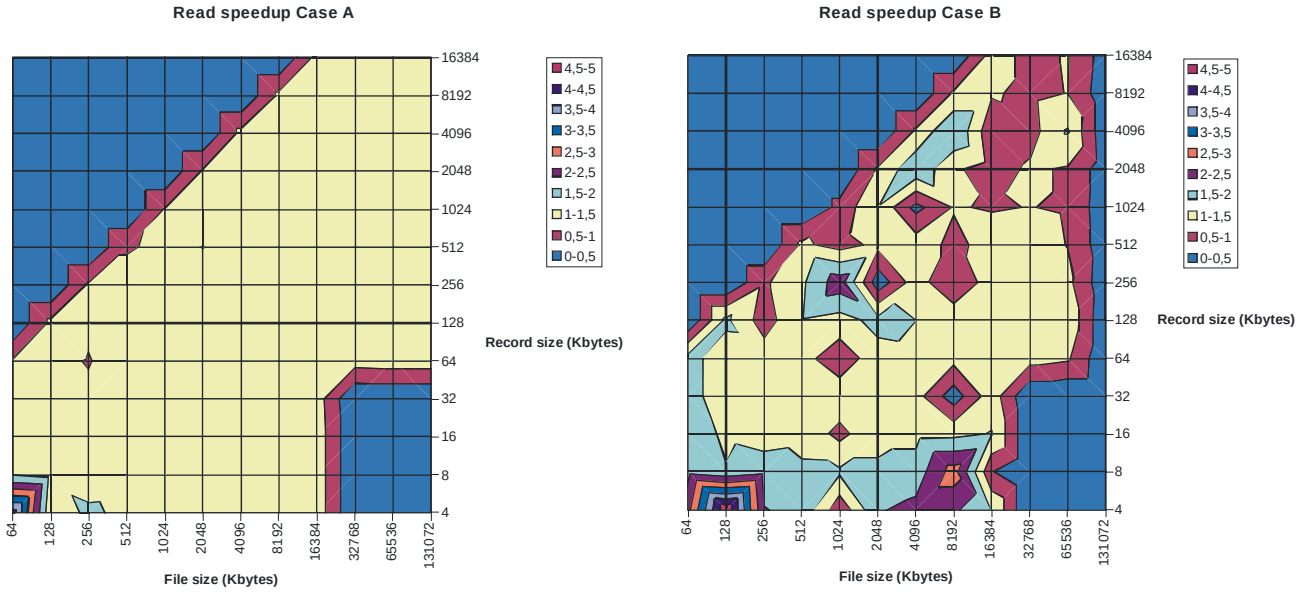

Figure 3. Read operation speedup comparison for case A (Windows XP) and case B (Linux).

Experiment 2

This experiment compares the behavior of the same Linux host OS when running over two base OS configurations.

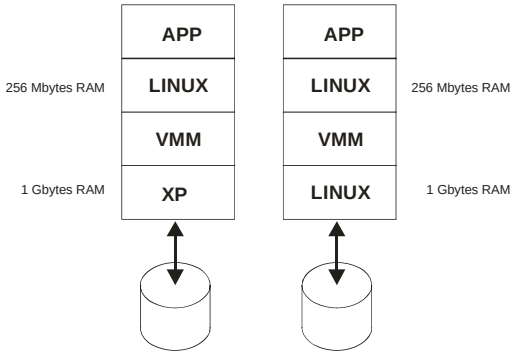

Figure 4. Operating environments to compare in experiment 2.

In the following figure we include the speedup on the *read* and *write* operations, obtained dividing the measures with Linux as base OS by the measures with Windows XP base OS. Values less than 1 indicates that with Windows XP performs better and values greater that 1, the Linux base OS performs better.

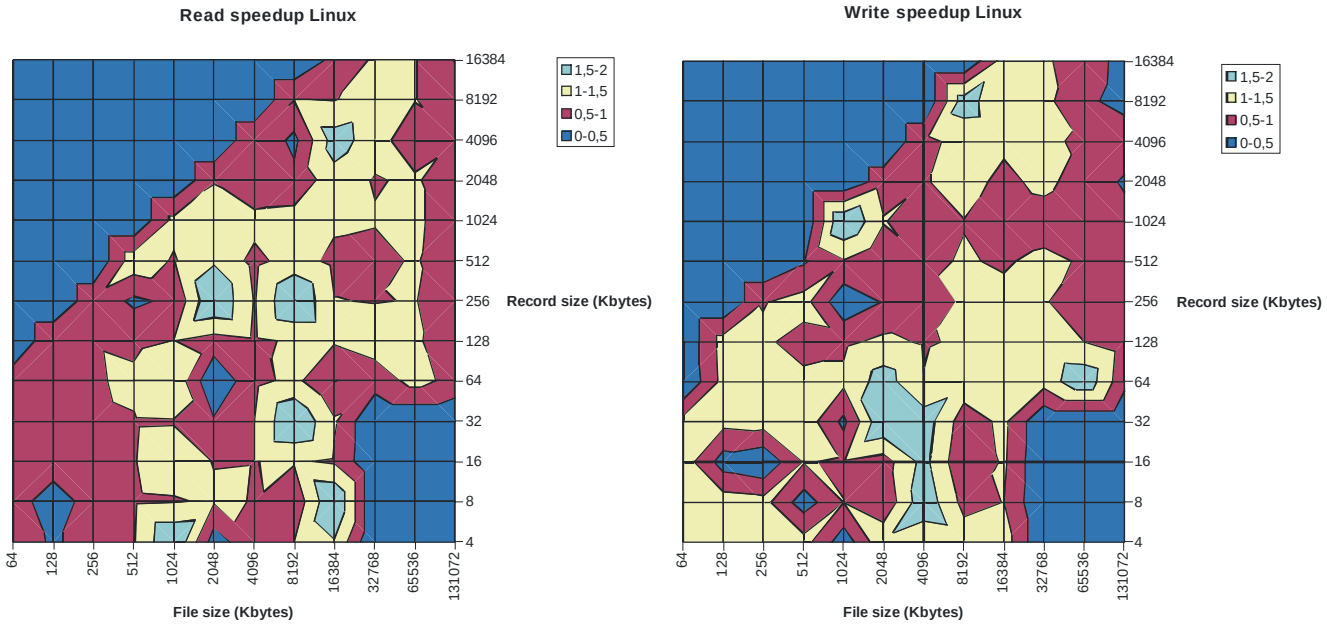

Figure 5. Read and write operations speedup comparison for experiment 2.

Observing the graphs depicted at Figure 5, we find that *read* operation has a generalized better performance for file sizes between 512 and 65536 Kbytes when running on Linux base OS. Results of *write* operation show that it has superior performance on Linux base OS.

Experiment 3

This experiment compares the behavior of the same Windows XP host OS running in two base OS configurations:

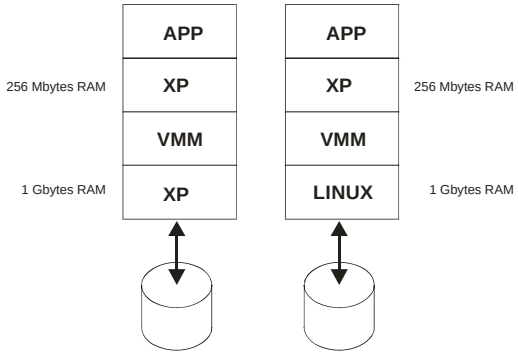

Figure 6. Operating environments to compare in experiment 3.

In Figure 7 we represent the speedup on the *read* and *write* operations, obtained dividing the measures with Linux as base OS by the measures with Windows XP base OS. Values less than 1 indicates that with Windows XP performs better and values greater that 1, the Linux base OS performs better.

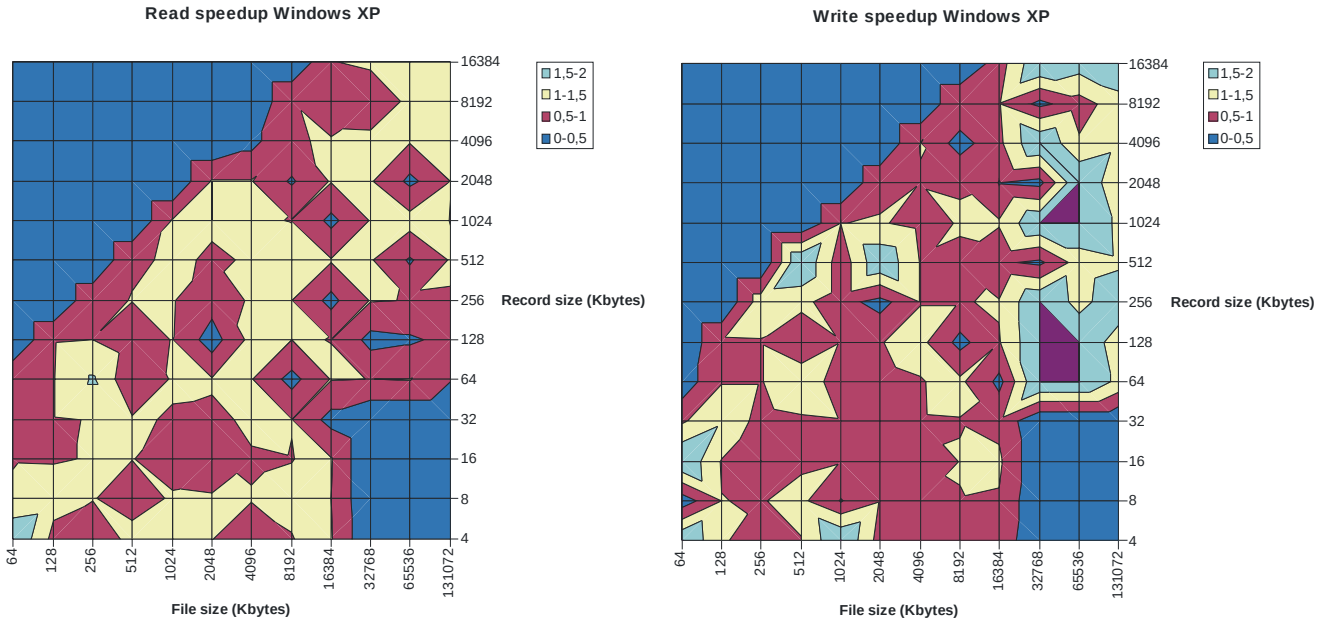

Figure 7. Read and write operations speedup comparison for experiment 3.

Results from Figure 7 indicate that the *read* operation does not have a well defined performance pattern. On the other hand, results of *write* operation show that it has higher performance on Linux base OS when file size is between 32768 and 131072 Kbytes.

3. CONCLUSIONS

We have presented a study on the overhead introduced by using a VMware Virtual Machine Monitor when managing a file system. The study has evaluated various operating environments to characterize the file system performance differences when an I/O benchmark is executed.

The VMM adds a layer of management between the base OS and the host OS that runs over it. Experiments results indicate that the overhead introduced by the VMM is generalized on the I/O operations. Depending on the operating environment, we have found that some write operations run faster over a VMware VM than running on the base OS.

When varying the base OS while the host OS is the same, experiments 2 and 3 show that Linux host OS behaves better over Linux as base OS.

4. REFERENCES

[1] VMWare
http://www.vmware.com/

[2] IOzone filesystem benchmark
http://www.iozone.org/