



Technical Note #12: Performance for High Volume File Systems

Technical Note	# 12. Performance for High Volume File Systems
Last Modified:	November 17, 2003
Version	1
Area:	Siebel File System on Windows Server 200X
Siebel Releases:	Up to 7.1
Windows Releases	Server 200X
SQL Server Releases	N/A

Summary:

Several Customers have had performance problems when their Siebel File Systems reaches 1.5 million files and more. Microsoft has done some internal testing copying large numbers of files into a single directory and has found that significant performance improvements can be made if customers disable 8.3 naming generation. It is on by default.

The Problem

This problem only affects the Siebel File System with very large numbers of files. The Siebel File System uses a non hierarchical directory structure with all files residing in a single folder. Microsoft recommends a hierarchical structure for file systems and Microsoft Tools such as explorer.exe and NTBackup.exe perform well in a hierarchical directory structure but these tools suffer performance problems with single folders that contain very large numbers of files. The NTFS file system itself is fully able to handle very large numbers of files but as the test below indicate, performance suffers when 8.3 file name generation is enabled.

During the testing done internally, 32K files were copied using the following naming scheme.

```
<Threadnum>_<batchnum>_<filenum>.dat
```

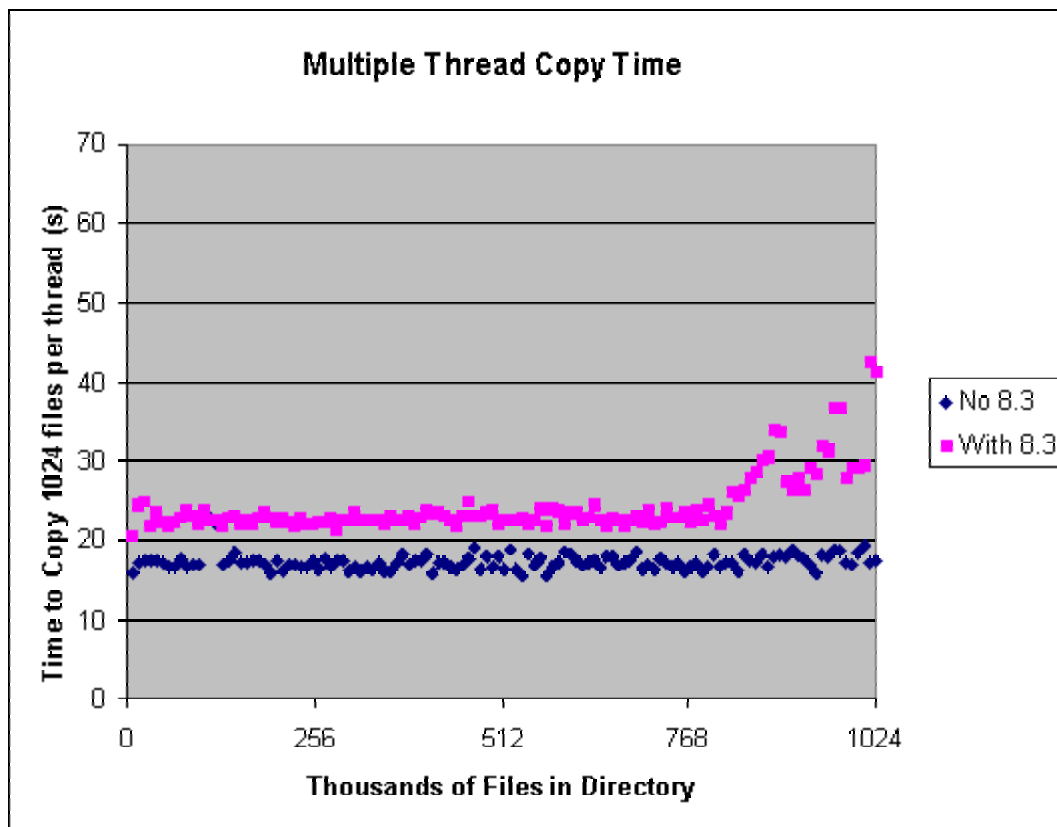
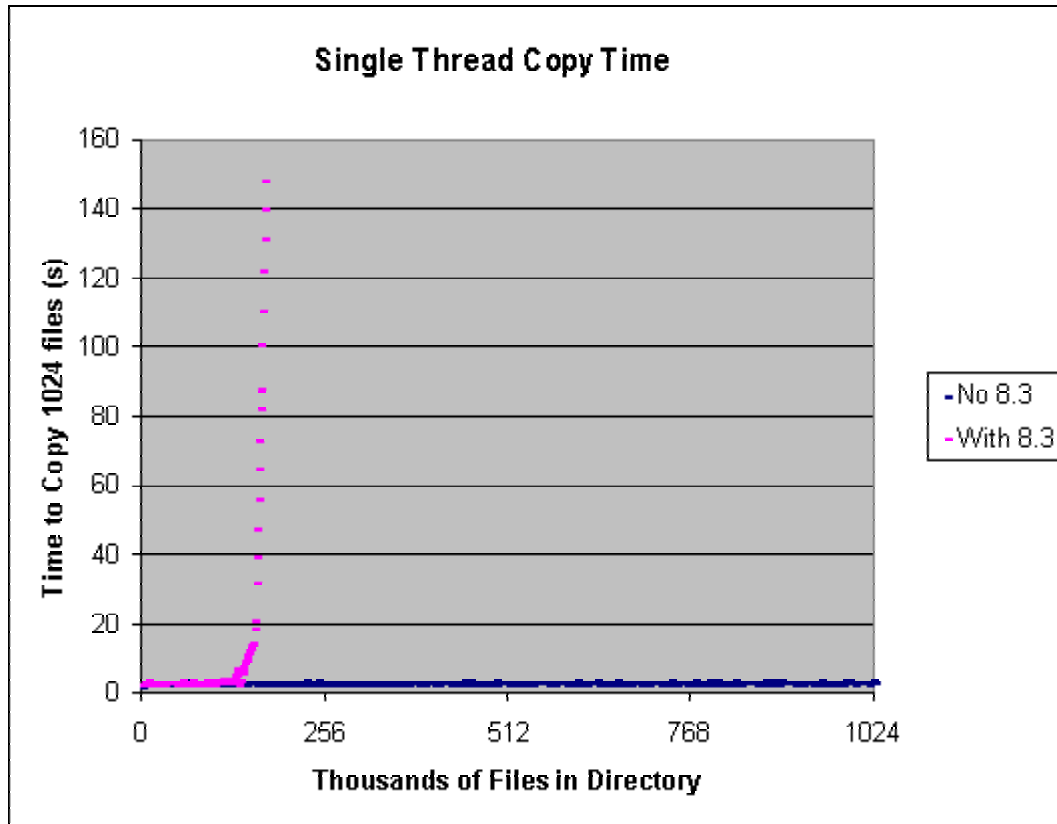
Where filenum is in the range from 1 to 1024. A million files were copied with Threadnum * batchnum * 1K equal to 1M.

Tests were performed on an 8 processor machine using both a single threaded work load and an 8 thread work load.

In the single threaded test, a few file names (25/1024) exceeded 8.3 during the 10th batch, but most files 925/1024 are longer than 8.3 after the 100th batch.

In the 8 thread case the files exceed 8.3 after the 80th batch and most files are short until the 800th batch.

Look at the charts below to see how file copying time takes off exponentially when most of the files are longer than 8.3.



Solution

Disable automatic short file name generation.

Regedt32 select:

```
"HKLM\system\CurrentControlSet\Control\FileSystem\NtfsDisable8dot3NameCreation"=1. (default is 0)
```

For an expanded discussion on how to disable automatic short file name generation, refer to Knowledge Base Article [210638 How to Disable Automatic Short File Name Generation](#)

Implications and Known Problems of disabling short file name generation.

The KB Article above (210638) specifically states that disabling short file name generation will cause existing files in short name format to become unavailable to 16 bit applications. They will however still be available to 32 bit applications.

[814844](#) PRB: Error Message: "Setup was unable to validate the logged user..."

<http://support.microsoft.com/?id=814844>

[811008](#) PRB: You Receive Error Message: "The replication agent is not registered"

<http://support.microsoft.com/?id=811008>

[290075](#) Distributed File System Links May Not Be Visible When You Connect to the

<http://support.microsoft.com/?id=290075>