Condusiv™ Technologies

PERFORMANCE EVALUATION GUIDE

I/O BOTTLENECKS - A WINDOW INTO YOUR FUTURE

Testing the effects of file fragmentation on your SAN's performance

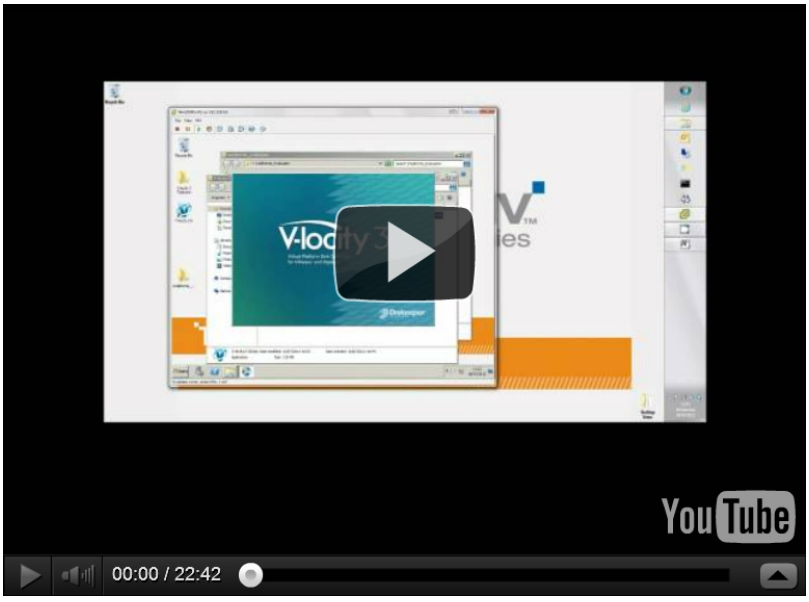

Executive Summary

Implementing the latest SAN technologies in either a physical or virtual environment produces a significant performance gain as a result of its consolidated, block level data storage. While things may be running smoothly now, at some point a SAN administrator will realise that the SAN storage is no longer performing as well as it once did.

The administrator may wonder what is causing all the unexpected I/O, and increasing the amount of storage available is usually the easiest solution. However, in many cases adding storage is unnecessary because space on the SAN is not the problem.

One of the most significant issues, and the most unrecognized, is fragmentation on top of a SAN storage system. With the implementation of a SAN, many Windows Server administrators believe that fragmentation, which they accepted and dealt with when using direct attached storage, has gone away. With this test we will show you that file fragmentation is still very much a key consideration to extending the life and performance of your SAN.

Fragmentation and SANs

The problem of file system fragmentation has been with us for decades. Direct-attached storage systems were plagued by it first, now SAN infrastructures are suffering its effects - inherited from the overlying file systems.

The effects of file fragmentation in SAN storage often manifest in reduced application performance and inefficient use of storage. Application response times begin to degrade, the time necessary to load large files and applications grows longer, and the overall user experience is negatively impacted.

End-users begin to feel that their computer is slowing down, leading to help desk calls with complaints about network performance or some other imagined problem. The reality is that file fragmentation is simply causing data manipulation times to increase to the point where the delay becomes perceptible to the end-user.

Understanding the Problem

"The reality is that although SAN vendors do an excellent job of optimising the performance and reliability of the networking component that they control (the storage), the SAN, by the nature of its design, is server operating system agnostic. This means that the SAN has no control over how the operating system treats its storage. The SAN's role is to provide the storage; the operating system's job is to deliver that data in a manner that works best."

– Windows IT Pro, Maximize the Performance of your Windows SAN Infrastructure

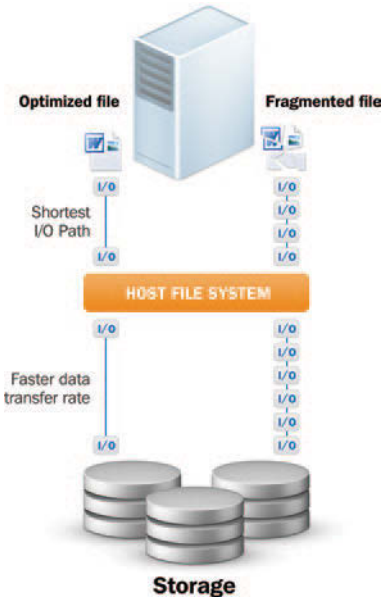

Fragmentation over SAN storage

The nature of a decoupled storage model precludes the operating system from being aware of the type of storage it's using. This greatly limits the operating system's ability to optimise for a particular storage model. To better visualize the problem, begin with the closed system of the Windows® operating system. Within a workstation or server, any I/O request has a minimum number of I/Os required based on file size. One could consider this the "Shortest I/O Path," with the greatest efficiency of data throughput.

File fragmentation generates excess I/O along the I/O path by increasing the number of searches required to complete an I/O operation. This directly results in slower data access as well as slower writes.

With the rise of abstracted storage layers, the problem of fragmentation is not miraculously eliminated – it is worsened. File system fragmentation also increases the amount of data accesses required to read/write at the block level due to how file pointing works between the OS and SAN. No matter how abstract you make the storage model, pieces of a file exist somewhere physically and each of those pieces are still associated, footstep by footstep, with the file representation within the operating system. I/O that occurs at the operating system level also occurs down at the SAN layer; excess fragmentation leads to more I/Os above as well as creating more I/Os below.

Solving the Problem

Condusiv I/O performance software offers technologies designed to help companies make their SANs run more efficiently and prevents performance degradation of their applications. While traditional methods of handling fragmentation may cause some extra write activity or thin-provisioning growth in SANs, Condusiv's preventative fragmentation technology – IntelliWrite, is fully compatible and extremely beneficial for SANs. With CogniSAN™ technology, Diskeeper® 12 Server detects external resource usage within a SAN allowing for increased performance without intruding in any way into SAN – layer operations. V-locity® virtual platform optimiser does the same with Virtual Machines.

Intelliwrite® fragmentation prevention technology

Historically, fragmentation has been addressed reactively, after it has already happened through the defragmentation process. In the "early days", fragmentation was addressed by transferring files to clean hard drives. Then manual defrag programs were introduced. The next step was scheduled defragmenters with varying degrees of automation. Truly automatic defrag was finally achieved with the development of InvisiTasking technology by Condusiv Technologies in 2007.

However, in spite of all the progress made with defrag methods, when fragmentation occurs, the system is wasting precious I/O resources by writing non-contiguous files to scattered free spaces across the disk and then secondly, using more I/O resources to defrag. Clearly the best strategy is to prevent the problem before it happens in the first place and always work with a clean, fast disk.

Based on knowledge of the Windows file system, Condusiv's patented IntelliWrite technology controls the file system operation and prevents most of the fragmentation that would otherwise occur.

Intelliwrite® file write optimisation technology is found exclusively in both Diskeeper® 12 storage performance software and V-locity 3 virtual platform optimiser.

> *"I tested Diskeeper on a 64bit Windows 7 system. My initial wariness against 3rd-party tools for critical tasks like disk defragmentation quickly vanished. The tool has obviously been developed to high quality standards and made a reliable impression on me. The user interface is intuitive and gives just the right amount of options; I liked the calculation and display of performance improvement estimates. Hyperfast – the optional flash drive module with support for special algorithms appropriate to SSDs – fits in seamlessly. My favorite is the "intelliWrite" option to configure Diskeeper such that disk fragmentation is prevented in the first place, rather than fixed in the aftermath.*
>
> *- Franz Kasparec, Cloud & Virtualization Business Development Manager EMEA-East, EMC Corporation*

Test Methodology

Testing IntelliWrite involves two primary evaluations: fragmentation and I/O throughput. Fragmentation evaluation is accomplished by observing the file write process (both at initial file creation and at file extension, or growth) and, at a granular level, inspection of fragmentation occurring. Given two mirror file creation/extension cases, one with IntelliWrite enabled should generate only 1-5% of the total fragmentation occurring in the other. Evaluation of I/O throughput & efficiency reveals the effect of IntelliWrite on the I/O generated on a system.

Testing described here involves use of a utility which generates a specified quantity of files, then simulates activity with those test files. The simulation creates a high level of file fragmentation. The IntelliWriteBench utility records both the number of fragments generated and the time taken to complete operations. Running the utility with and without IntelliWrite enabled yields comparative fragmentation as well as file writes time results. Additionally, observation of I/O shows the benefits of IntelliWrite in reducing the amount of work required by the system in file writes.

Testing should be performed with limited or no other system activity and on a non-system volume in order to further diminish any other contributing factors to performance results. The end result should be a true apples-to-apples comparison.

Notes: The test volume should have at least 4GB of available free space, and Performance Monitor logging must be enabled. The evaluation runs at a normal user priority; since it is running an accelerated simulation, a great deal of I/O activity will be generated. The system will likely be very slow during this test—this is intended, in order to truly evaluate how IntelliWrite performs under intense load. Depending on available system resources and evaluation mode selected, the evaluation may take 10 minutes or many hours to complete.

Test Utilities

IntelliWriteBench.exe – Creates numerous files, and then modifies the test files with additional information. This is intended to simulate activity on a system at an accelerated rate through the writing and increase of file sizes. These modifications cause significant file fragmentation in the absence of IntelliWrite.

The Source Code is available for viewing at the following location:
http://downloads.condusiv.com/pdf/IntelliWriteBench_Source_Code.pdf

Contig.exe – CLI (Command Line Interface) single-file defragmenter and file creator from Microsoft® / Sysinternals. This is employed by IntelliWriteBench to analyse the test files for fragmentation.

Diskeeper 12 or V-locity 3 – Storage performance software for Windows workstations and servers.

Windows Performance Monitor – System performance monitoring utility within Windows.

IntelliWrite_Evaluation.bat – Batch file which automates the usage of IntelliWriteBench, Contig, & Diskeeper 12 or V-locity 3 in testing while also creating and storing all relevant Windows Performance Monitor data during testing.

Test Steps

NOTE: Please read carefully through all Test Steps before proceeding—there are important notes at the end about viewing evaluation results.

1. Download IntelliWrite_Evaluation.zip from the following location : http://downloads.condusiv.com/zip/Intelliwrite_Evaluation.zip

2. Extract and save IntelliWrite_Evaluation.bat, IntelliWrite_Evaluation_Results.xlsx, and IntelliWriteBench.exe to the same location on your intended test system.

3. Download Contig.zip from the following location: http://technet.microsoft.com/en-us/sysinternals/bb897428 Extract and save Contig.exe into the same folder as the IntelliWrite Evaluation files.

4. Install Diskeeper 12 - http://downloads.condusiv.com/trialware/Diskeeper12-Server-45day.exe or V-locity 3 - http://downloads.diskeeper.com/trialware/V-locity3eo9kb-30day.exe

   a. Open Diskeeper 12 or V-locity 3 and ensure that IntelliWrite is enabled for the intended test volume within the Volume Pane. Also turn off Auto Defrag and Space Reclamation features. Only Intelliwrite should be enabled otherwise I/O's caused by defragmenting or zeroing could skew the results.

5. Perform the evaluation using IntelliWrite_Evaluation.bat.

   a. Single right-click on IntelliWrite_Evaluation.bat within the previously selected test directory, then select 'Run as Administrator'.

   b. Select the evaluation mode (Brief, Standard, or Extensive).

      I. Brief Mode generates 45 files per test run (out of 6), and requires approximately 2GB of available free space on the test volume. Total run time should be approximately 5-10 minutes.

      II. Standard Mode generates 150 files per test run (out of 6), and requires approximately 6GB of available free space on the test volume. Each test run will take between 5-10 minutes to complete—approximately 1 hour total.

   c. Extensive Mode generates 750 files per test run (out of 6), and requires approximately 29GB of available free space on the test volume. Each test run will take several hours to complete.

   d. Select whether the Evaluation is being performed on SAN storage.

      I. If this option is selected, 10 worker threads are employed during each test run, rather than the standard 3. This is done in order to more closely resemble work load on a SAN storage system. Note that this means that the number of files being generated is slightly more than 3x the standard values—as is the required available free space. A Brief Mode Evaluation performed with the SAN option will take approximately 15 to 30 minutes to complete.

e. Enter the test volume letter when requested; the evaluation uses this information to generate a working directory named 'IntelliWrite_Evaluation_Temp' for storing performance information. This working directory is removed automatically after the test has concluded. Wait while the evaluation occurs; IntelliWrite_Evaluation.bat handles all calls to IntelliWriteBench.exe, Contig.exe, as well as creates and records all PerfMon counters.

f. If this is the first time that Contig.exe has run on this machine, you will be prompted to accept a license agreement. Click 'Accept'.

g. Wait while the evaluation occurs; IntelliWrite_Evaluation.bat handles all calls to IntelliWriteBench.exe, Contig.exe, as well as creates and records all PerfMon counters.

h. After testing has concluded, you will be asked whether Microsoft Excel® is installed on the test system. If so, the test procedure will automatically launch test results from C:\IntelliWrite_Evaluation_Results\IntelliWrite_Evaluation_Results.xlsx; in order to view all performance data from the testing graphed & charted, please click Enable Content in order to dynamically update the results file, then click save.

i. Notes about viewing results:

I. If the test machine is not equipped for viewing Excel documents, all results will be copied back to the initial test directory in a subfolder named IntelliWrite_Evaluation_Results. Please copy the entire IntelliWrite_Evaluation_Results directory to the C:\ volume on the machine where you intend to view results and then launch IntelliWrite_Evaluation_Results.xlsx. Ensure that you Enable Content within Excel.

II. If you plan to rerun this evaluation for multiple volumes, you should rename the IntelliWrite_Evaluation_Results directory and add a volume letter to the name to distinguish them.

III. If you Enable Content again on an IntelliWrite_Evaluation_Results.xlsx file from a previous test run, it will overwrite its content with newer test results—only Enable Content the first time you open a Results file, then save it.

Additional Reading

Inside IntelliWrite Technology:

http://downloads.condusiv.com/pdf/Inside_Diskeeper_12_with_IntelliWrite.pdf

Maximize the Performance of Your Windows SAN Infrastructure (Windows IT Pro paper):

http://downloads.condusiv.com/pdf/Performance-Windows-SAN.pdf

Best Practices for using Diskeeper storage performance software on Storage Area Networks (SANs):

http://downloads.condusiv.com/pdf/Best_Practices_for_using_Diskeeper_on_SANs.pdf

Condusiv Technologies Europe
2nd Floor, Garland Court,
Garland Road, East Grinstead,
West Sussex, RH19 1DN
+44 (0)1342-821-300
www.condusiv.co.uk

Condusiv Technologies
7590 N. Glenoaks Blvd.
Burbank, CA 91504
800-829-6468
www.condusiv.com