

GPFS Scans 10 Billion Files in 43 Minutes

**Richard Freitas, Joseph Slember, Wayne Sawdon and
Lawrence Chiu**

**IBM Advanced Storage Laboratory
IBM Almaden Research Center
San Jose, CA 95120**

© Copyright IBM Corporation 2011.

Trademarks

IBM and other IBM products and services are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries or both.

Any other trademarks, registered trademarks, company, product or service names may be trademarks, registered trademarks or service marks of others.

Table of Contents

1	Introduction	1
2	Active management of large data repositories	3
2.1	Application demands continue to grow	3
2.2	Active Management	4
2.3	Scaling	5
2.4	GPFS.....	5
2.5	Disk Technology Trends	6
2.6	A conundrum	9
2.7	Solid-state storage	10
3	Demonstration.....	11
3.1	Test stand.....	12
3.2	GPFS requirements for solid-state storage	13
3.3	Procedure	16
3.4	Results	17
4	Discussion	21
5	Glossary	23
6	References	24

List of Figures

FIGURE 1: SYSTEM PERFORMANCE TREND.....	3
FIGURE 2: AVAILABLE STORAGE CAPACITY TREND	4
FIGURE 3: AREAL DENSITY TREND.....	6
FIGURE 4: MAXIMUM SUSTAINED BANDWIDTH TREND.....	7
FIGURE 5: AVERAGE LATENCY TREND	8
FIGURE 6: AVERAGE SEEK TIME TREND.....	9
FIGURE 7: TEST STAND BLOCK DIAGRAM.....	11
FIGURE 8: TEST STAND PHOTOGRAPH.....	13
FIGURE 9: VIOLIN PERFORMANCE FOR 16 KB BLOCK	14
FIGURE 10: VIOLIN PERFORMANCE FOR 256KB BLOCKS	14
FIGURE 11: VIOLIN SYSTEM READ LATENCY DENSITY	15
FIGURE 12: COMPLEMENTARY CUMULATIVE LATENCY DISTRIBUTION	16
FIGURE 13: AGGREGATE READ OPERATIONS PER SECOND	18
FIGURE 14: AGGREGATE WRITE OPERATIONS PER SECOND	19
FIGURE 15: AGGREGATE READ BANDWIDTH.....	19
FIGURE 16: AGGREAGE WRITE BANDWIDTH.....	20
FIGURE 17: AGGREGATE CPU UTILIZATION.....	20
FIGURE 18: FILE SYSTEM SCALING	21

Abstract

By using a small cluster of ten IBM xSeries® servers, IBM's cluster file system (GPFS™), and by placing file system metadata on a new solid-state storage appliance from Violin Memory, IBM Research demonstrated, for the first time, the ability to do policy-guided storage management (daily tasks such as file selection for backup, migration, etc.) for a 10-billion-file environment in 43 minutes. This new record shatters previous record by factor of 37. GPFS also set the previous record in 2007.

The information processing power consumed by leading business, government and scientific organizations continues to grow at a phenomenal rate (90% CAGR). This growth is one of the factors driving the growth in data repositories. The amount of online digital data should exceed 1800 EB by the end of 2011 and continue to grow at a rate 40-60% per year [1]. This explosive growth of data, transactions and digitally aware devices is straining IT infrastructure and regular data management operations. The task of managing storage: backing up, migrating to appropriate performance tiers, replication and distribution is overburdening this infrastructure. It is not possible with existing solutions to manage 10 billion files actively today.

Unfortunately, the performance of the current, commonly-used storage device -- the disk drive -- is not keeping pace with the rate of performance growth needed by business and HPC systems. Recent advances in solid-state storage technology deliver significant performance improvement and performance density improvement, which is suitable for future storage systems matched to the needs of a growing information environment.

This document describes a demonstration that shows GPFS taking 43 minutes to process the 6.5 TBs of metadata needed for a file system containing 10 Billion files. This accomplishment combines the use of enhanced algorithms in GPFS with the use of solid-state storage as the GPFS metadata store. IBM Research once again breaks the barrier of GPFS scalability to scale out to an unprecedented file system size and enable much larger data environments to be unified on a single platform and dramatically reduce and simplify the data management tasks, such as data placement, aging, backup and replication.

1 Introduction

In our increasingly instrumented, interconnected and intelligent world, companies struggle to deal with an increasing amount of data. This paper illustrates a ground breaking storage management technology that scanned an active ten billion-file file system in 43 minutes, shattering the previous record by a factor of 37. The IBM General Parallel File System™ (GPFS) [2] architecture, invented by IBM Research-Almaden, represents a major advance of scaling for performance and capacity in

storage, more than doubling the performance seen in the industry over the past few years.

According to a 2008 IDC study [1] there will be 1800 EB of digital data in 2011 and it will be growing 60% per year. This explosive growth of data, transactions and digitally aware devices is straining IT infrastructure, while storage budgets shrink and user demands continue to grow. At current levels, existing solutions can barely cope with the task of managing storage: backing up, migrating to appropriate performance tiers, replication and distribution. Many such file systems go without the kind of daily backup that industry experts would expect of a large data store. They will not be able to manage when file systems scale to 10 billion files.

Such growth places businesses under tremendous pressure to turn data into actionable insights quickly, but they grapple with how to manage and store it all for their current set of applications. As new applications emerge in industries from financial services to healthcare, traditional systems will be unable to process data on this scale, leaving organizations exposed to critical data loss.

Anticipating these storage challenges decades ago, researchers from IBM Research – Almaden created GPFS, a highly scalable, clustered parallel file system. Already deployed in enterprise environments with one billion files to perform essential tasks such as file backup and data archiving, this technology's unique approach overcomes the key challenge in managing unprecedented large file systems with the combination of multi-system parallelization and fast access to file system metadata stored on a solid-state storage appliance.

GPFS advanced algorithms make possible the full use of all processor cores on all of these machines in all phases of the task (data read, sorting, and rules evaluation). GPFS exploits the excellent random performance and high data transfer rates of the 6.5 TB solid-state metadata storage. The solid-state appliances sustainably perform hundreds of millions of IO operations, while GPFS continuously identifies, selects and sorts the right set of files from the 10 billion-file file system. Performing this selection in 43 minutes was achieved by using GPFS running on a cluster of ten 8-core systems and four Violin Memory solid-state memory arrays.

In this document, we will describe the current environment for active management of large-scale data repositories and describe a demonstration of GPFS coupled with the use of high performance solid-state storage appliances to achieve this breakthrough.

2 Active management of large data repositories

2.1 Application demands continue to grow

The information processing power consumed by leading business, government and scientific organizations continues to grow at a phenomenal rate. Figure 1 shows the historic growth in supercomputer system performance from 1995 until 2010, with projections to 2015. The growth rate is roughly 90% CAGR [3]. Further, a comparison of the computing power of a current system, such as the IBM System x® iDataPlex® dx360 [4], with the computing power of a system from 2000, shows that they are in rough parity. One can infer that general enterprise computing needs will be met by technology that was available in the TOP500 leader roughly ten years earlier.




Figure 1: System performance trend

Associated with the growth in system overall computational performance is a growth in the size and performance of the system's data repository. The growth in size is driven by both the increase in computational capability and the continual increase in information to store. IDC has reported that the growth rate in digital information is 60% CAGR, with the current size of the world's digital information at 1800 EB. This is depicted in the chart [1] shown in Figure 2. More importantly, the combination of the scale of the computers combined with much larger data set sizes supported by parallel file systems allows new classes of questions to be answered with computer technology.

“When dealing with massive amounts of data to get deeper levels of business insight, systems need the right mix of data at the right time to operate at full speed. GPFS achieves high levels of performance by making it possible to read and write data in parallel, distributed across multiple disks or servers.”⁵

2.2 Active Management

Active management is a term used to describe the core data management tasks needed to maintain a data repository. It includes tasks such as backup migration, backup, archival, indexing, tagging for files, etc. A feature these tasks share is the need to enumerate or identify a file or set of files in the data repositories for later processing. The process of managing data from where it is placed when it is created, to where it moves in the storage hierarchy based on management parameters, to where it is copied for disaster recovery or document retention, to its eventual archival or deletion is often referred to as information life-cycle management, or ILM. Such tasks need to be done regularly – even daily.

Information Created vs Available Storage Capacity




Figure 2: Available storage capacity trend

However, performing such tasks is disruptive to normal operation or normal operation precludes performing active management tasks in a reliable consistent manner. This has forced these ILM

tasks to be performed when production work is not being done. Typically, this means that the tasks are only performed on the graveyard shift and must be completed before normal daytime workload arrives.

2.3 Scaling

As the performance of the world's leading supercomputers and business systems has increased, their data repositories have become enormous and extremely complex to manage. However, the amount of time available to perform the necessary management tasks has not changed. Therefore, the speed of active management must increase.

One solution has been to split the data repository and spread the pieces across multiple file systems. While this is possible, it burdens the application designer and the user unnecessarily. The burden placed on the application designer is one of prescience. A priori, he must think of all the possible configurations and cope with the needed data being in one of many file systems. The user's burden is to allocate data statically to the file systems in a manner that will allow the application to operate effectively. Changes in performance needs, available capacity or the approach to disaster recovery may necessitate reallocating the data. Doing this "on-the-fly" in a statically allocated data repository is a very error-prone process.

2.4 GPFS

The solution used by IBM's General Parallel File System (GPFS) is to "scale out" the file system to meet the user's needs. Capacity or performance can be increased incrementally by adding new hardware to an existing file system, while it is actively used. GPFS provides concurrent access to multiple storage devices, fulfilling a key requirement of powerful business analytics and scientific computing applications that analyze vast quantities, of often, unstructured information, which may include video, audio, books, transactions, reports and presentations. GPFS scales to the largest clusters that have been built, and is used on some of the most powerful supercomputers in the world. The information lifecycle management function in GPFS acts like a database query engine to identify files of interest. All of GPFS runs in parallel and scales out as additional resources are added. Once the files of interest are identified, the GPFS data management function uses parallel access to move/backup/archive the user data. GPFS tightly integrates the policy driven data management functionality into the file system. This high-performance engine allows GPFS to support policy-based file operations on billions of files. GPFS is a mainstay of technical computing that is increasingly finding its way into the data center and enhancing financial analytics, retail operations and other commercial applications.

There is a fifteen-year record of success behind GPFS. For large repositories, this is an important attribute. Critical applications using large data repositories rely on the stability of the file system. The cost of rebuilding it can be substantial. GPFS is designed to remain on-line and available 24x7, allowing both hardware and software changes without interruption.

2.5 Disk Technology Trends




Figure 3: Areal density trend

Disk capacity has improved immensely over the last fifty-seven years. This is exemplified by the chart of areal density vs. year of introduction shown in Figure 3. The growth rate of areal density has varied from 25% to 100% over this period and is now at 25–40%. The maximum capacity of disk drives has tracked this growth in areal density through changes in form factors, disk diameters, media changes, head technologies, etc. It is unlikely to ever be 100% CAGR again, but the current rate will likely be supported for the next few years.

Unfortunately, the performance of the disk drive is not keeping pace with the rate of performance improvement shown by business and HPC systems. The maximum sustained bandwidth (MB/s) and transaction rate (IO/s) have followed a different path. The bandwidth is roughly proportional to the linear density. So, if the growth in linear density and track density were equal, then one would expect the growth rate for linear density to be the square root of the areal density. That would make it about 20% CAGR. But, if you examine the recent history of maximum sustained disk bandwidth that is shown in Figure 4, you will see that it is more likely to fall within the range of 10

- 15%. Generally, the track density has grown more quickly than the linear density. Currently a high performance disk drive would have a maximum sustained bandwidth of approximately 171 MB/s [6]. The actual average bandwidth would depend on the workload and the location of data on the surface. Further, current projections do not show much change in this over the next few years.




Figure 4: Maximum sustained bandwidth trend

The disk transaction rate has an even more complicated story. The transaction rate is the inverse of the average access time. The access time is the sum of two components: the average disk-latency time and the average disk-seek time.



Figure 5: Average latency trend

The average disk-latency is $\frac{1}{2}$ the rotational time of the disk drive. As you can see from its recent history, shown in Figure 5, the latency has settled down to three values 2, 3 and 4.1 milliseconds. These are $\frac{1}{2}$ the inverses of 15000, 10000 and 7200 revolutions per minute, respectively. It is unlikely that there will be a disk rotational speed increase in the near future. In fact, the 15000 RPM drive and perhaps the 10000 RPM drive may disappear from the marketplace. A slower drive, such as one at 5400 RPM, may appear in the enterprise space. If this reorganization of the disk rotational speed menu does occur, it will have been driven by the successful combination of solid-state storage and slower disk drives into storage systems that provide the same or better performance, cost and power.




Figure 6: Average seek time trend

The recent history of disk seek time is shown in Figure 6. The seek time is due to the mechanical motion of the head when it is moved from one track to another. It is improving by about 5% CAGR. In general, this is a mature technology and is not likely to change dramatically in the future. It seems to be trending toward values at around 2 ms for high performance disk drivers.

The combination of slow growth in seek time and a near stand-still in rotational latency or even possible increase in average rotational latency, lead to the conclusion that the average access time for enterprise disks is not likely to be better than 3–4 milliseconds and that commercial disk drives (7200 RPM) are not likely to be better than 6–8 milliseconds. Access times greater than this are quite possible.

2.6 A conundrum

The growth in the size of data repositories is expected to continue at 60% CAGR. This forces an increase in the size of the metadata that needs to be processed to complete active management tasks. Since these tasks must be completed overnight (typically), the scan rate for the metadata must increase. Therefore, the performance of the metadata storage system must increase. This situation coupled with the very modest forecast in the growth rate for disk drive performance (~ 5–10% CAGR), present the industry with a conundrum. To continue the growth in the performance of the metadata storage system, the number of disk drives in the system must grow at a matching rate even though the capacity of the drives far exceeds that needed to store the metadata. Fortunately, there is an alternative. There is a new class of storage, solid-state storage, which may ameliorate this situation.

2.7 Solid-state storage

NAND flash devices are one of several new kinds of storage subsystems. Typical flash devices or chips currently have a read access times of 20–50 us and write times of approximately 2 ms [7]. Hundreds or even thousands of NAND flash chips combined with a controller that is designed to work with flash yield a high performance solid-state storage subsystem. The controller is sophisticated and handles buffering, virtual addressing and the special control needs of flash. As a result, a solid-state storage appliance can be built that has a read access time of 90 us and a write access time of 20 us [8]. The subsystem form factor may be that of a disk drive, an adapter card or a larger storage appliance. Flash is not the only possible solid-state storage technology that may enable the design and manufacture of such fast storage systems. A new class of storage technologies, SCM (Storage Class Memory), is under intensive development in the industry [9]. If flash falters, then one of these is likely to take over.

3 Demonstration

This document describes a demonstration that shows GPFS taking 43 minutes to process the 6.5 TBs of metadata needed for a file system containing 10 Billion files. That is roughly 37 faster than the demonstration of the same activities on a 1 Billion-file system performed in 2007 [5]. A combination of new algorithms in GPFS and the use of a solid-state storage appliance for storing the metadata were used to achieve these results.




Figure 7: Test stand block diagram

In this section, we will describe the hardware used in the demonstration test stand, the procedure used and the results obtained.

The demonstration took place in a laboratory at the IBM Almaden Research Center in San Jose, CA. IBM provided the ten servers and the IB switch and Violin Memory provided the four Violin Solid-state Memory Arrays.

3.1 Test stand

Figure 7 shows the block diagram for the test setup. The components in the system are:

1. IBM System x3650 M2 servers
 - Ten used
 - CPU: 2.8 GHz dual quad core x86_64
 - Processor cache size 12 MB
 - 32 GB of DRAM
2. Violin Memory 3205 Solid-state Storage Systems
 - Four used
 - Aggregate total raw capacity of 10 TB
 - Aggregate bandwidth 5 GB/s
 - 1.8 TBs formatted per 3205, aggregate usable capacity 7.2 TBs
 - Two 14x 128GB partitions
 - Two 10x 180 GB partitions
 - Aggregate 4 KB read operation rate > 1 MIOPS
 - Typical write latency at 4KB: 20us
 - Typical read latency at 4 KB: 90us
3. SilverStorm Infiniband switch
 - Model 9024
 - 24 port
 - 10/20 Gb/s

Figure 8 shows a photograph of the actual physical setup used for these experiments.




Figure 8: Test stand photograph

3.2 GPFS requirements for solid-state storage

We chose to use Violin Memory solid-state storage array for these experiments because of its high performance, robustness and density.

The Violin memory array in the 3000 series systems offers very high, sustained read and write performance. Figure 9 and Figure 10 below show the measured transaction rate of a single Violin 3000 for a range of read percentages from 0 to 100%, for a range of queue depths from 1 to 300 and for representative block sizes of 16KB and 256KB. The figures show that the system performs well against a mixture of reads and writes at both block sizes. In the demonstration, four such boxes are used. Their aggregate transaction rate capability at 16 KB with a mixture of reads and writes is more than 300 KIOPS and their aggregate read bandwidth at 256 KB blocks is almost 5 GB/s. (See section below for the discussion below about the performance needed for the demonstration)




Figure 9: Violin Memory performance for 16 KB block

The latency profile for the Violin 3000 is excellent. Figure 11 shows the latency density function for 16 KB transfers and Figure 12 shows the complementary cumulative latency distribution response.




Figure 10: Violin performance for 256KB blocks

The Violin 3000 has a very sharp latency peak at roughly 90us for 16 KB reads as is shown in Figure 11. Figure 12 shows that 99% of the reads complete within 100 us and 99.999 complete in less than 500us. Together these show that the system has a fast and very predictable latency.




Figure 11: Violin read latency density

Lower latency I/Os mean faster task completion. Consider two systems, A and B, one using solid-state storage with an access time of 100 us and the other using conventional disk drives with an access time of 5 ms. If you have a task to complete that consists of 100,000 storage operations and they are all parallelizable, then 50 drives will match the performance of the solid-state storage. However, if 1% of the I/O operations must be serial, then 100 drives are needed for matching and if 2% must be serial, then it is impossible for a disk-based system to match the solid-state system.

The Violin system is a robust storage device. It gains its robustness from three factors. First, each individual data record is encoded to protect against errors, etc. Second, data are stored across multiple flash cards as a RAID3 array. Parity is calculated on-the-fly and stored at the same time the data are stored. There is no RAID5 write penalty. In addition, when data are read, the whole stripe is read, including the parity and the parity of the stripe is checked immediately. This provides added protection. Finally, the Violin system is twin-tailed. This allows to overall system to maintain access to the data by “failing-over” if a link or server fails.

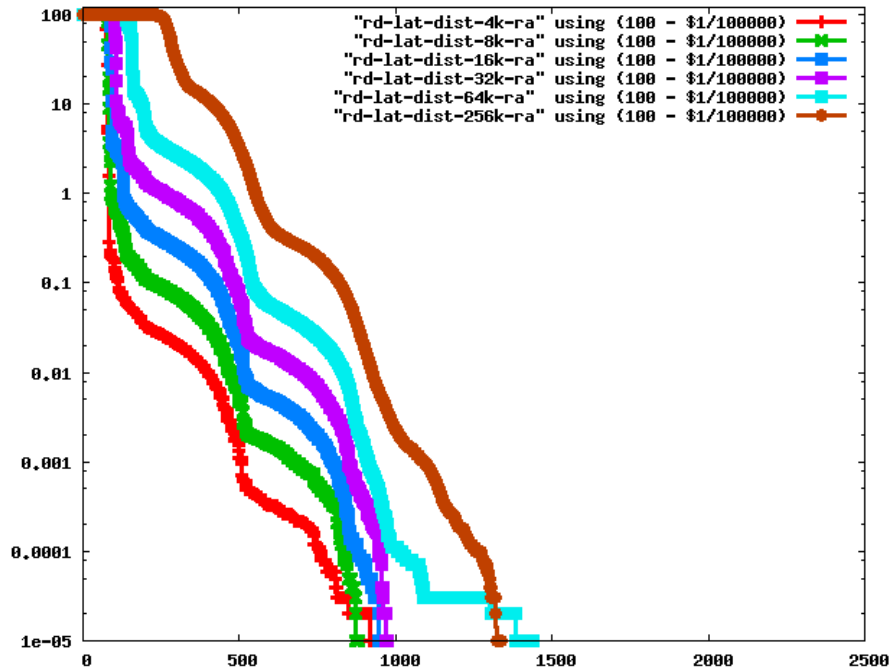


Figure 12: Complementary cumulative latency distribution

Another important advantage that solid-state storage has over hard-disk storage in applications requiring very high performance is volumetric density. In situations, such as the one discussed in this white paper, where the performance requirement is high, but the capacity requirement is not, a system such as Violin's can have a significant space advantage. 6.5 TB of formatted storage is all that is needed to meet the storage requirements for the metadata. However, it needs 4 GB/s in read performance from that storage. Four current SATA drives would be sufficient to house the metadata, but meeting the bandwidth requirements would entail the use of 200 drives. Currently that would be 12U for Violin memories against at least 40U (plus space for the disk controllers) for disk drives.

3.3 Procedure

We used the product GPFS, version 3.4, with enhancements that have subsequently been released into the product service stream. The Violin storage is configured into 48 LUNs and used to create a standard file system using a 1 MB block size with all data and metadata residing solely in solid-state storage. We then populated the file system with 10 Billion zero length files. Since the file system policy scan being measured accesses only the files' metadata, we omitted all file data. Alternatively, we could have added storage to the file system and stored the file data on-line into either solid-state or conventional hard disk or even off-line to tape. In all cases, the location of the data has no impact on policy scan's performance. The files were spread evenly across a little more than 10 Million directories. The file system was unmounted and remounted between each

benchmark run to eliminate any cache effects. The policy scan was invoked across the 10 Billion files as we measured the time required to complete the “scan”. For each of the computers and devices used in the demonstration, we collected measurements for the following metrics: read bandwidth, read transaction rate, write bandwidth, write transaction rate and the CPU utilization of the client/NSD computers. These metrics were sampled every second throughout the “scan” and the aggregate results are shown below in section 3.4 .

The policy scan runs in two distinct phases: a rapid directory traversal to collect the full path to each file followed by a sequential scan over each file’s inode to collect the remaining file attributes. This second phase also performs the policy rule evaluation, which selects the files that match the policy rules. For this benchmark, we used a single policy rule that matched none of the files.

The directory traversal phase consists of a parallel traversal of the file system directory structure. A master node assigns work to multiple processes running on each node. Each process is assigned a disjoint sub-tree of the directory structure, beginning with the root directory (or an arbitrary directory or set of directories within the file system). In GPFS, each directory entry contains a “type” field to identify sub-directories without requiring a stat on each file. Thus the directory traversal requires all directories to be read exactly once. This required over 20 Million small reads operations – one to read each directory’s inode and one to read the directory’s data block. When this phase is complete, the full path to every file and that file’s inode number is stored in a series of temporary files. The number of temporary files created depends on the number of nodes used and the number of files in the file system. The files are kept relatively small to optimize parallel execution in the second phase.

The sequential scan phase begins by assigning the temporary files to the processes running on each node. Each set of temporary files contains a range of inodes in the file system. The files within the range are sorted into inode number order and the files’ inodes are read directly from disk via a GPFS API, thus providing the files’ attributes such as owner, file size, atime and it includes the files’ extended attributes, which are simple name-value pairs. Since the files’ inodes are read in order, we perform sequential full block reads with prefetching on the inode file. This test required reading over 5 TB of inode data. Each files’ inode is merged with its full path, then submitted to the policy rule evaluation engine which checks each policy rule against the files’ attributes checking for a match. Files that match a rule are written to a temporary file for later execution or may be executed immediately, before the scan completes.

3.4 Results

The full GPFS policy scan over 10 Billion required 43 minutes. As shown in the graphs below, the directory traversal phase completed in 20 minutes. The directory traversal is IOP intensive,

composed of small reads at random locations in the file system. This phase of the scan greatly benefits from the underlying Violin storage, but our workload did not approach the IOP capability offered by the 4 Violin memories. The first phase also wrote approximately 600 GB bytes of data spread across more than 5000 temporary files. The five charts below clearly show the performance characteristics during this phase – a high IO operation requirement, with relatively little data read, but a moderate amount of data written. The CPU requirements during this phase are primarily due to process dispatching for the jobs reading directories.

The second phase for the policy scan reads the temporary files created by the first phase, sorts them into inode order, then reads each file’s inode and performs the policy rule evaluation. For 10 Billion files, this phase required about 23 minutes and read 5.5 TB of data. As shown in the graphs below this phase is bandwidth bound, even when spreading the flash memory from 2 fully-loaded Violin memories across 4 boxes making each half full. Also note that the CPU cost for sorting the files into inode order and the cost for policy rule evaluation is masked by the wait times for reading the data.




Figure 13: Aggregate read operations per second




Figure 14: Aggregate write operations per second




Figure 15: Aggregate read bandwidth




Figure 16: Aggreage write bandwidth




Figure 17: Aggregate CPU utilization

4 Discussion

Figure 18 shows the scan rate in millions of files per second that GPFS has measured and announced, plotted against the size of the file system being managed. Today's demonstration at 10 Billion files is highlighted as A and the graph projected on to 100 Billion files. A system size that is likely to be needed within a few years.




Figure 18: File system scaling

While four disk drives would probably be sufficient to store the meta-data for 10 Billion files, they are not nearly fast enough to process it. Processing the metadata using disk drives would require at least 200 disk drives to achieve the same processing rate. Since disk drive performance is growing slowly; managing the metadata for 100 Billion files would require approximately 2000 disk drives. In 2007, the 1 Billion file scan required around 20 drives. This progression will have significant impact on the data center housing future large data repositories. A standard storage drawer is all that is required to house 20 drives. For 200 drives, the space requirement moves up to a standard data center rack assuming 2 1/2" disk drives. For 2000 drives, the space required would be 5 to 10 data center racks. The slow growth in the performance of disk drives will entail a significant increase in the number of drives dedicated to metadata processing and that will increase

the space, power and number of controllers needed. In addition, since the areal density is growing much faster than the performance, the storage efficiency of the drives dedicated to metadata processing will be poor. Clearly, there is a need for a different medium, solid-state devices, that have the performance characteristics necessary to solve these problems.

This demonstration is important for enterprise customers facing an explosion in data. The increase in performance of disk drives has not matched the increases in capacity, nor is it expected to in the near future. This trend increases the cost of large file systems and will limit their ability to grow. Fortunately, the performance and capacity offered by today's solid-state devices meets the requirements for today's largest file system and will meet the needs for the next generation as well. Building a file system to serve 100 Billion files will require half a rack of solid-state storage for the metadata, compared to 5-10 full racks of disk drives to achieve the same level of performance. Using solid-state storage significantly reduces the power required to operate the devices, the power required for cooling, and the cost of the additional space itself.

Today's demonstration of GPFS' s data management showed its ability to scale out to an unprecedented file system size and enable much larger data environments to be unified on a single platform. Using a combined platform will dramatically reduce and simplify the data management tasks, such as data placement, aging, backup and replication. This reduced cost of operation, plus the reduced cost for power, cooling and space, is critical for continued data growth. This scalability will pave the way for new products that address the challenges of a rapidly growing multi-zettabyte world.

IBM Research has and continues to build on a long history of innovative storage technologies helping clients not only manage data proliferation, but harness data to create new services. In the past year alone, IBM storage products included over five significant storage innovations invented by IBM Research including IBM Easy Tier®, Storwize® v7000, IBM Scale Out Network Attached Storage™ (SONAS), IBM Information Archive™, and the IBM Long Term File System™ (LTFS).¹⁰

5 Glossary

atime	Most recent inode access time
CAGR	Compound Annual Growth Rate
EB	Exabyte -- 10^{18} bytes
GPFS	General Parallel File System
HPC	High Performance Computing
ILM	Information Life Management
inode	An internal data structure that describes individual file characteristics such as file size and file ownership
MIOPS	Millions of I/O operations per second
NAND	NAND Flash -- Flash memory chip where the basic cell is based on a nand (negative and) gate
NSD	Network Shared Disk
SSD	Solid-state Disk
TCO	Total Cost of Ownership
ZB	Zettabytes -- 10^{21} bytes

6 References

- [1] Gantz, John F., etal. “The Diverse and Exploding Digital Universe”. IDC, March 2008, Framingham, MA.
- [2] <http://www-03.ibm.com/systems/software/gpfs/>
- [3] <http://www.top500.org/>
- [4] <http://www-03.ibm.com/systems/x/>
- [5] <http://www-03.ibm.com/press/us/en/pressrelease/22405.wss>
- [6] <http://www.seagate.com/staticfiles/support/disc/manuals/enterprise/cheetah/15K.6/SAS/100466193b.pdf>
- [7] <http://www.micron.com>
- [8] <http://www.violin-memory.com>
- [9] Freitas, Richard and Chiu, Lawrence. “Solid State Storage: Technology, Design and Systems”. Usenix FAST2010, www.usenix.org/events/fast10/tutorials/T2.pdf.
- [10] <http://www-03.ibm.com/systems/storage/>