

Umbrella File System: Storage Management Across Heterogeneous Devices

JOHN A. GARRISON and A.L. NARASIMHA REDDY
Texas A&M University

With the advent of and recent developments in Flash storage, device characteristic diversity is becoming both more prevalent and more distinct. In this paper, we describe the Umbrella File System (UmbrellaFS), a stackable file system designed to provide flexibility in matching diversity of file access characteristics to diversity of device characteristics through a user or system administrator specified policy. We present the design and results from a prototype implementation of UmbrellaFS on both Linux 2.4 and 2.6. The results show that UmbrellaFS has little overhead for most file system operations while providing an ability better to utilize the differences in Flash and traditional hard drives. With appropriate use of rules, we have shown improvements of up to 44% in certain situations.

Categories and Subject Descriptors: D.4.3 [Operating Systems]: File Systems Management;

General Terms: Design, Measurement, Performance

Additional Key Words and Phrases: device characteristics, flash drives, policy-driven storage, namespaces

ACM File Format:

GARRISON, JOHN A. and REDDY, A.L. NARASIMHA. 2009. Umbrella File System: Storage Management Across Heterogeneous Devices. *ACM Trans. Storage* ?, ?, Article ? (? 2009), ?? pages.

1. INTRODUCTION

Novel storage devices based on flash memory are becoming available with price/performance characteristics different from traditional magnetic disks. Storage systems in the future will likely incorporate both flash-memory based devices and magnetic disks. It is necessary for storage systems and file systems to be designed to exploit the unique characteristics of these different devices. While current storage systems exploit devices of different characteristics, some of the characteristics of flash drives such as limited write cycles warrant us to revisit this problem.

With the emergence of solid-state flash RAM drives, the differences between different types of devices are becoming more pronounced in some characteristics. Flash drive capacities are approaching the 100 GB range, and their read and write speeds are approaching those of more traditional magnetic disks [Samsung 2007]. These drives have much faster seek times than their magnetic counterparts, due to their lack of moving parts, but often have a limited number of times that the drive can be written. Flash drives also typically have more uniform performance depending on file size, where magnetic disks typically perform better with larger file sizes, as can be seen in Section 4.

Various methods have been proposed to compensate for and exploit this diversity in device characteristics. Most storage systems use memory for caching and manage memory and storage devices together. Systems such as HP's AUTORAID [Wilkes et al.

Authors' address: 214 Zachry, Texas A&M University, College Station, TX 77843; email: {jgarrison, reddy}@ece.tamu.edu.

Permission to make digital/hard copy of part of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date of appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee. Permission may be requested from the Publications Dept., ACM, Inc., 2 Penn Plaza, New York, NY 11201-0701, USA, fax: +1 (212) 869-0481, permission@acm.org
© 2001 ACM 1530-0226/07/0900-ART9 \$5.00 DOI 10.1145/1290002.1290003 <http://doi.acm.org/10.1145/1290002.1290003>

1996] have been designed to work at the device level by automatically moving blocks between different devices. Flash memory is being used as a cache to improve disk drive performance [Trainor 2006]. These methods work at the device level, and thus present a single view of the device to the file system and do not expose the diversity of the device characteristics to the file systems and applications above the file system.

While this transparency of device characteristics may reduce the challenges in dealing with devices of different characteristics, it also prohibits flexibility in exploiting these differences to the advantage of individual applications. For example, if a user deems it ideal to store many small read-only files on a flash drive while storing the remaining files on a magnetic disk in a storage system, currently the file systems and storage systems do not provide the user any mechanisms to accommodate such a system organization. While Object Storage Devices (OSDs) [Project T10 2004] can potentially provide such flexibility, OSDs require extensive changes to storage systems, file systems and clients. High level hints have been shown to be valuable in improving I/O performance [Patterson et al. 1992].

Like devices, applications also have remarkable diversity. While some applications utilize primarily sequential storage accesses and thus would not benefit from caching, some other applications have much more locality in their access patterns and would benefit greatly from caching. Some files may be read-only while others may be frequently updated.

Current file systems map files to devices based on the namespace organizations. Typically, a file system's data is mapped to a fixed set of devices based on the way the namespace is mapped to the underlying devices. If a user wants to locate some of his files (say /usr/foo) on one device (a flash drive) while locating other files in /usr on another device, most current file systems do not provide a simple mechanism to allow this since all of the files under /usr are mapped to a fixed set of devices in the storage system. This mapping is determined by the relation between the namespace and devices in the storage system.

In general, current file systems do not provide applications above them a choice of where to locate files or which devices to employ. This choice is only available at the entire file system granularity and not at an individual file level.

In this paper, we propose a solution to allow diversity of applications to be matched to the diversity of devices at an individual file level. The matching of application characteristics to device characteristics can be done by a system wide policy or based on the user's preferences. We propose UmbrellaFS to make this "policy-based storage" possible. We describe the details of UmbrellaFS and our implementation of policy-based storage below.

File systems interact with both applications and the underlying storage systems, and their location at the intersection of these two systems with diverse characteristics makes them uniquely suited to take advantage of both device and application characteristics. While the file system does not have large amounts of information about the underlying device where it resides, the system's administrator typically is aware of device characteristics and differences. UmbrellaFS allows the device characteristics to be exposed at the granularity of a file system. Hence, matching application characteristics to device characteristics translates into placing files of different characteristics onto appropriate file systems. UmbrellaFS provides a mechanism for files in a single user directory to be located appropriately on multiple file systems or devices through a user directed policy.

As an example configuration that can benefit from UmbrellaFS, consider a system with three different types of storage, traditional magnetic storage devices, a smaller

capacity flash based device, and another magnetic disk with hardware encryption support. There are various file systems [Halcrow 2005] [Hohmann 2007] [Gough 2006] and even device level [Seagate 2007] encryption options available to users. Consider a situation where some user files need to be encrypted on such a system. If we employ an encryption file system encompassing all the devices, we don't harness the hardware encryption support available at one of the devices (or encrypt some files twice). Encrypting all the files may impose overheads that may be undesirable for multimedia files and other files that do not have major security requirements. If we employ normal file systems on all the drives and rely on the device-level encryption, the user has to carefully place the sensitive files on this device to ensure their protection.

UmbrellaFS can be employed in this scenario to direct files to the appropriate underlying storage device while allowing the files to be viewed in the same directory. This avoids the overhead of encrypting everything, as well as allowing the user to map files that best utilize the encrypting device for storing sensitive documents and the flash drive for storing read-only files (for example).

It has been observed that current device level coding schemes utilized to protect bit errors on the drive surface may not guarantee 100% data protection as drive capacities get larger. While data files require 100% bit accuracy, multimedia files can tolerate some bit errors without suffering any significant loss of quality. It is possible that in the future drives with different levels of bit error protection may become available. Again, such diversity in devices can be easily mapped to the applications through UmbrellaFS developed in this paper.

Opportunities like this and the others a user could conceive serve as motivation for UmbrellaFS. With UmbrellaFS, users can select file systems that serve the particular characteristics of the data sets without compromising the user's convenience of organizing files any way he or she sees fit. Different files can be stored in separate file systems while being presented to the user in a single directory. The user can deal with a variety of files without worrying about the particular file system where a file is stored. With the UmbrellaFS rules framework, we give users the option of setting up the mapping from files to underlying storage in a manner that optimizes file usage.

In Section 2 we outline our approach to this problem, and how UmbrellaFS is set up. Section 3 details our prototype, and Section 4 presents benchmark results. A number of example scenarios that show the utility of UmbrellaFS are shown in Section 5. Section 6 presents related work and Section 7 concludes and presents directions for future research.

2. OUR APPROACH

2.1 Hourglass Model

The Virtual File System (VFS) [Kleiman 1986] maintains a one to one mapping between the namespace and the underlying file systems. If one is in the directory "/user1", the files "/user1/file1" and "/user1/file2" are not located on separate file systems. This can be seen in Figure 1. This one to one mapping ties the user's perception of the system to the underlying device characteristics. If this one to one mapping could be broken, such that one directory in the namespace could contain files from many underlying file systems, then this could enable exposure of device characteristics to individual files and allow new opportunities and options for the user.

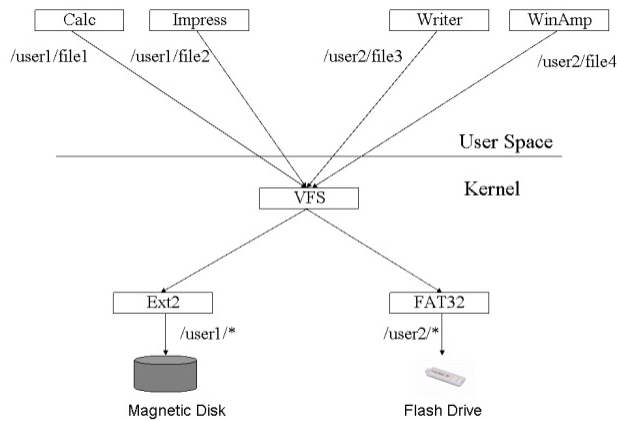


Fig. 1. Hourglass Model with traditional Virtual File System

By inserting another layer into the kernel, below VFS and above the underlying file systems, it is possible to place files from multiple file systems in the same namespace. This layer would look and behave like the underlying file systems to VFS, and look and behave like VFS to the underlying file systems. When VFS receives a command, it would pass it into this new layer, which would then redirect the command to the appropriate underlying file system. This method would allow both VFS and the underlying file system to continue to operate normally, and would not require any modification of those parts of the kernel. While VFS continues to provide the important function of demultiplexing among the many different underlying file system implementations, the underlying layer, which we call UmbrellaFS, provides the function of unifying the namespaces and demultiplexing the namespace operations among the different underlying file systems.

UmbrellaFS lies below VFS and interfaces between VFS and the underlying file systems. UmbrellaFS provides a unified name space for all the underlying file systems while also allowing policy based placement of files onto those file systems. The proposed approach with UmbrellaFS is shown in Figure 2. While VFS provides an hourglass model for file system implementations with a single interface, UmbrellaFS provides a similar function for the namespace and individual users' files, as seen in Figure 2.

The Virtual File System provides a common API for applications to interact with file systems. This allows the file system's actual implementation of various functions such as writing, reading, etc. to be hidden from the application. For example, when an application accesses file /user1/file1 in Figure 1, VFS receives the system call and then calls the appropriate function for the underlying file system. When the underlying function returns, VFS returns that value up to the application. With UmbrellaFS, while the underlying mechanisms are slightly different, from the user's perspective the file access proceeds as normal. UmbrellaFS appears as a single file system to VFS, and consequently to the user. When the user accesses the file, VFS calls the appropriate UmbrellaFS call. UmbrellaFS then translates the filename to the correct underlying file system and calls the appropriate function for that file system. When the file system function returns, UmbrellaFS returns the value to VFS as if it were the base file system on which the file resides, and VFS returns to the application.

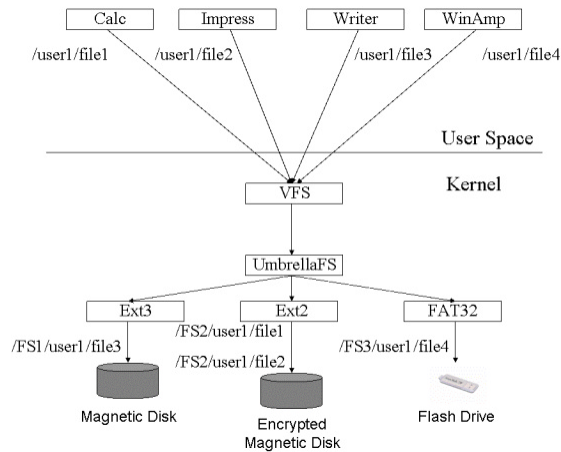


Fig. 2. Hourglass Model with UmbrellaFS

In this manner, UmbrellaFS can leverage the disparity in device characteristics by mapping particular files to particular underlying file systems which appropriately exploit the device characteristics on which they reside. Additionally, UmbrellaFS increases the opportunity for future file system development with respect to particular underlying devices. Rather than requiring a full change of file systems, incremental development and installation of underlying file systems is permitted. UmbrellaFS exposes the device characteristics to the user or system administrator at the level of native file systems. For example, a traditional file system might be employed on a magnetic disk drive, while a log based file system is employed on a flash drive in order to conserve energy and lengthen the drive lifetime [Mathur et al. 2006]. Based on user-controlled policies, files can be mapped to the underlying device by mapping individual files (/user1/file1, /user1/file2... in Figure 2) to the native file systems (/FS1, /FS2... in Figure 2).

As an example of UmbrellaFS's utility, take content specific storage. If the types of files are known, they can be mapped to an underlying storage device that best meets their typical access patterns. For example, UmbrellaFS can direct files to underlying file systems based on the file suffix. Multimedia files with extensions such as .jpg, .gif, and .bmp files could be stored in /images. Text files such as .txt and .doc extensions could be stored in /text, and video files such as .avi files could be stored in /video. The user is presented with a single directory structure with all of the files present. For example, /user/dir1/foo.txt and /user/dir1/foo.jpg would be stored in /text/dir1/foo.txt and /images/dir1/foo.jpg respectively. With files mapped in this manner, it would be possible to put text files on a file system with redundancy, while the various multimedia files could be stored on a system with wide striping for faster access, etc. Figure 3 shows the user's view of the namespace, as well as how the files are stored in the underlying storage.

There are alternatives to the implementation of a new file system approach as we describe in the paper. It is possible for users to attempt to exploit device and application level diversity manually. A user can direct files to appropriate file systems on top of storage devices that best work with those files' access patterns and the applications that use them. This approach has a number of drawbacks, including a great deal of time and thought required by the user for uniform implementation. With the manual user-based approach, file locations are still tied to the underlying storage device, and large amounts of user intervention are required for nearly every operation.

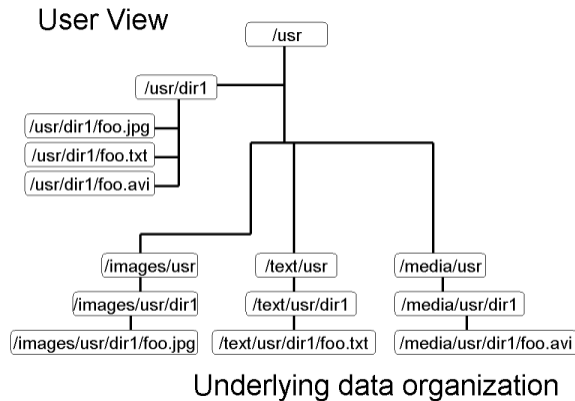


Fig. 3. Example mapping of user namespace based on file contents.

Additionally, it would be possible that rather than providing a combined namespace as UmbrellaFS does one could rely on soft links to achieve a similar effect. The process could be automated, and files could be placed on appropriate underlying devices while links to the files are displayed to the user in a combined directory. In addition to the lack of transparency in this method, it is likely that the overhead and work involved in maintaining soft links and an additional combined directory structure would meet or exceed that encountered by UmbrellaFS's current implementation.

UmbrellaFS is designed to provide flexibility in storage allocation. We present here some examples of how this flexibility can be exploited by the user. Much of this flexibility is derived by breaking the one-to-one link from the user namespace to the underlying storage devices. The flexibility is envisioned to be exploited through a policy module administered by the user or system administrator, the details of which are described below.

2.2 System Design

UmbrellaFS functions as a stackable file system [Zadok et al. 2006], residing below VFS and above the underlying file systems. All user interactions with the file systems pass through UmbrellaFS and are directed to the appropriate underlying file system. Some operations, such as open, remove, read, and write which operate on a single file are directed only to the underlying file system on which the file resides. Operations which list the contents of directories such as *ls* and *du* must be mapped to all the underlying file systems, and the results must be collated and passed back to the calling operation.

It is possible to provide different functionality by treating the failures on the branches differently when operations such as *ls* and *du* are executed on multiple branches or file systems. In one implementation, any failure on any branch is treated as a failure of the entire operation (an OR operation on the failures). In a second implementation, data from successful branches can be passed up, suppressing the failures on some branches (an AND operation on the failures). While more general functions can be employed in reporting failures, these two cases are considered in this paper, in two separate applications of UmbrellaFS.

In order to deal with the differences in directories (which may or may not reside on multiple underlying file systems) and files (which should only reside on a single underlying file system), UmbrellaFS maintains additional data about each file. This data

includes values which indicate the branch where the file currently resides as well as the branches where a directory might exist.

The key element of UmbrellaFS is the policy decision module. When any file system operation is invoked by an application, the subsequent system call sends control to the UmbrellaFS. UmbrellaFS then evaluates the file based on the rules provided by the user at file system mount time, and sends the command to the appropriate underlying file system. UmbrellaFS stores these rules in UmbrellaFS's super block. Return values from the underlying file systems are likewise directed back to VFS, which then returns them to the application.

The policy module employs rules or policies that determine the allocation of files on different devices. For example, a policy can state that any file larger than 100 kB should be directed to a file system (say /FSRAID) on a RAID5 [Patterson et al. 1988] device. A second policy can decide that Dave's files be stored on a second device. These rules can be based on any attributes of the file, including namespace, size, access privileges, time of creation, etc. UmbrellaFS directs file system operations to underlying file systems based on the different criteria specified in the policy module. Both namespace rules as well as metadata rules based on inode values can be specified. The rules are evaluated in a two pass system. Metadata rules are evaluated in order, returning the first positive match to the list. The namespace rules reside in the list of metadata rules, and are evaluated when the namespace rule is reached in the list. The different namespace rules are evaluated based on length of match between the file in question and the rule, much like longest prefix matching in routing. In this case however, the namespace rule with the longest overall match, not just the longest prefix is returned and used to determine which underlying file system to use. Should there be additional metadata rules underneath the namespace rules and if there were no matches in the namespace rules, then once the namespace rules return the system continues to evaluate the rest of the rules. Should there be no match whatsoever, the system directs the operation to a default file system.

At time of file creation, much of the metadata is either unavailable or liable to change shortly into a file's lifetime. Values such as create time can be handled by using the current time of the operation, but values like file size can be more problematic. As a large file is being copied, it will initially be a small file. If a rule is set up based on file size, this could result in moving a file between underlying file systems soon after creation. In many cases this would be mitigated by the fact that the file that is moved will be less than the full file, however there is an impact that is examined in Section 4.

Similarly, at file creation, a file would need to be writable, even though it would ultimately be a read-only file. Putting read-only files on a particular file system could therefore involve an overhead of moving files across file systems after file permissions are changed. These overheads can be minimized by employing policies on group and global permissions. Alternately, the mostly-read characteristics can be assessed by other characteristics such as file suffixes ending in .out .exe. While complex rule structures may be needed to adequately describe and define the desired behavior, UmbrellaFS makes these types of rule structures possible.

The policy decision module is in many ways similar to a router in a network, directing traffic to the appropriate destination. Many of the policy matching rules (first match, longest name match etc.) resemble the order in which routing decisions are applied when multiple entries in a routing table match an incoming packet. In UmbrellaFS, the file is an analog of a packet and the destination file system is an analog of nexthop. While we have used the rules of first match and longest name match, other appropriate rules may be employed in other designs.

Figure 4 shows an example system with one metadata and three filename policy rules. These rules are designed to represent the situation described in the introduction of a system with three different types of storage, traditional magnetic storage devices, a smaller capacity flash based device, and another magnetic disk with hardware encryption support. The user, Dave, wants to map most of his personal files to the encrypted storage, read only files to the flash storage, and his video files and all other files to the magnetic storage.

Rule 1 is based on access type and dictates that all read only files should be located on /fs2, the flash drive. For the purposes of this example, a process such as the previously described planned transition of file permissions is assumed. Rule 2 is based on filenames and is elaborated in the second table. The filename rules state that all of user Dave's files should be placed on the encrypted device, /fs3. The exception is .avi files which should be placed preferably on /fs1 (the magnetic disk) first and then on /fs3. Finally, all other files should be directed to /fs1.

As can be seen from this example, a number of policies may apply to a given file. The specified rules are evaluated in order of specified priority (first rule 1, then rule 2, and then on to other rules if they are present). For example, if user Dave creates a read only .avi file in his home directory, it will be placed on /fs2 since rule 1 is applied before the filename rules are examined. A decision tree or other appropriate data structure can be used to reason about the effect of the policies on the storage of files on different devices. When evaluating the rules and determining which storage device should hold a file, it is possible that the characteristics that mapped a file to a particular device might change. For example, a file size based rule might be in effect, and a file has grown (or reduced) to the point that the rule structure indicates the file should be on a different file system. These situations are evaluated as they arise, and should a file reside in a file system that is different from the file system that the policy decision module decides it should reside in, that particular file is moved to the appropriate underlying file system. This evaluation is made primarily during file closing (with some exceptions noted below). This movement between native file systems (or underlying devices) may result in performance overheads, which we evaluate later.

Rule No.	Rule Type	Compare	Value	File system
1	Access type	=	Read only	/fs2, /fs1
2	Filename			

Comparison String	File System
/home/dave/*	/fs3
/home/dave/*.avi	/fs1, /fs3
/*	/fs1

Fig. 4. Same UmbrellaFS rules

Such lazy enforcement of policy on file closing cannot be useful with certain rules based on time of access. For example, if we want a file to be stored on a slow device if it hasn't been accessed in five days, we would need another mechanism to enforce this

policy. A periodic policy enforcer may be needed in such cases. This policy enforcer could simply traverse the directory structure, opening and closing each file. This would allow UmbrellaFS to make the appropriate choices for file relocation. The overhead from this type of operation could be minimized by scheduling the policy enforcer to run during off-peak usage times. There is no need for a specific program to move files in the underlying file systems, as this could lead to discrepancies between UmbrellaFS's understanding of the underlying file systems and their actual state. UmbrellaFS depends on the lower level file systems remaining in the state in which it left them. Attempting to circumvent UmbrellaFS and accessing the lower level file system directly could result in UmbrellaFS not finding underlying files because it is directing operations to file systems that do not contain the files anymore. In general, while UmbrellaFS can be circumvented by accessing the underlying file systems, this should always be avoided to prevent situations like that laid out above. This risk can be diminished by mounting a "dummy" directory over the underlying directory, effectively hiding the directory from the user.

It is possible that the administrator specified policy may not be well suited for the underlying storage systems. For example, the administrator may specify that video files be stored on a 100GB device. If more than 100GB of video files need to be stored in the system, the policy needs to be flexible to use the space on remaining devices that may not be so well suited to video files. In order to allow for these exception cases, the policy can specify a set of choices for locating a given type of files. The policy module tries to place files in the specified preferred order of devices. Hence, in some exception cases, a file system operation such as a read may involve multiple underlying file systems. We study the impact of this overhead in our experiments.

In these exception cases the writes will be mapped to the next choice of system. If a file were to grow to the point that it could not fit, then the file would have to be moved to the secondary storage system before additional writes are done. This writing is done on the fly as individual actions to the underlying file systems return errors.

In addition to making backup locations like this possible, UmbrellaFS can also explicitly forbid them. If, for example, the rule in question put sensitive files onto an encrypted drive, one would not want these files to be written to another location if that drive became full. Correct behavior in this case would be to return with an error indicating there is no more room on the device. Whether a backup location is allowed for a given rule is up to the user. In the example above, rules 1 and 2.2 provided backup systems while rules 2.1 and 2.3 only allowed one choice in locating the files.

We realize that inferring beneficial policy rules automatically based on file access behavior in guiding policy decisions may be feasible. We leave this to future work. Our current framework will be able to accommodate such an approach through the employment of additional metadata characteristics or file tags.

3. IMPLEMENTATION

UmbrellaFS is implemented as a loadable module for both the 2.6 and 2.4 Linux kernels. We leveraged Unionfs 1.0.14 [Wright et al. 2004] in the development of UmbrellaFS. Unionfs's ability to combine underlying file systems and representing them in a single directory is an integral part of UmbrellaFS's functionality; however modification of the underlying functions of Unionfs was required in order to accomplish the specific goals of UmbrellaFS. There are no substantial differences in the 2.4 and 2.6 implementations.

The modifications were primarily located in the rule evaluation framework, and the methods for importing the rules into the system at module load time. Some additional minor changes were made in many functions to convert from linear searches of directories to the particular file system targeting used by UmbrellaFS. No kernel code external to UmbrellaFS was modified.

Unionfs combines the contents of underlying file systems into one unified view for the user. UmbrellaFS expands upon this idea by providing the ability to direct files to particular underlying file systems using either namespace rules or other metadata. In addition, UmbrellaFS is designed to separate files into the appropriate underlying storage devices as they are written, instead of combining existing file systems into one. Much of Unionfs's overhead is circumvented in UmbrellaFS because rather than searching through all of the possible directories linearly each time when accessing a file, UmbrellaFS uses the metadata to direct the operation to the appropriate file system. The inherent overhead from additional system calls is present in both UmbrellaFS and Unionfs.

4. EVALUATION RESULTS

Benchmark tests were run on a Dell Optiplex GX620 with an Intel 3.2 GHz Pentium D processor and 2 GB of RAM. Tests were run using a Red Hat Fedora Core 3 Linux 2.6.9 kernel as well as a Debian 2.4.27 kernel. For tests the kernel only used 250 MB of RAM in order to prevent caching from skewing the results. The system disk was a Samsung 7200 RPM 250 GB SATA hard drive, while the tests were run on Seagate Cheetah 10k.7 73 GB SCSI hard drives. The SCSI hard drives were connected to the system via Adaptec 29320A SCSI controllers. Ext2 was used as the file system on all hard drives. The lower level file systems were remounted between each test to clear file system caches.

A variety of macrobenchmarks and microbenchmarks were chosen to explore the characteristics of UmbrellaFS. In addition to the benchmarks examined in [Traeger et al. 2008] such as Bonnie++ [Coker 2001], Postmark [Katcher 1997], and compilation of source code, particular microbenchmarks were developed in order to examine some of the potential issues in UmbrellaFS. Each test was run at least 3 times, and the averages are provided. Ninety-five % confidence intervals were calculated using the Student-*t* distribution. The half-widths of the intervals are less than 5% of the mean in most cases, and exceptions are noted in their particular sections.

UmbrellaFS's overhead was evaluated in a number of different situations. Raw device throughput, CPU and I/O intensive loads were examined. All tests were performed with exclusively inode rules as well as exclusively filename rules. When inode based rules were used, the worst case scenario (no match on the rules until the final rule) was used. Similarly, with the filename rules, rules were chosen which had at least some match to the files used so that as much time as possible would be spent comparing various string matches.

Basic overhead was examined when UmbrellaFS was mounted over a single underlying file system. Additionally, overhead was examined when UmbrellaFS resided over multiple underlying file systems, both when data was confined to a single branch, and when data was distributed between branches. The multiple branch results are not included, but do not differ significantly from the results shown.

4.1 Bonnie++

Bonnie++ is a benchmark which tests sequential accesses to a single file, or multiple files if the working size is larger than 1 GB. A working size of 2 GB was used for Bonnie++, so Bonnie wrote two 1 GB files. The results of the Bonnie++ benchmark are shown in Figures 5.

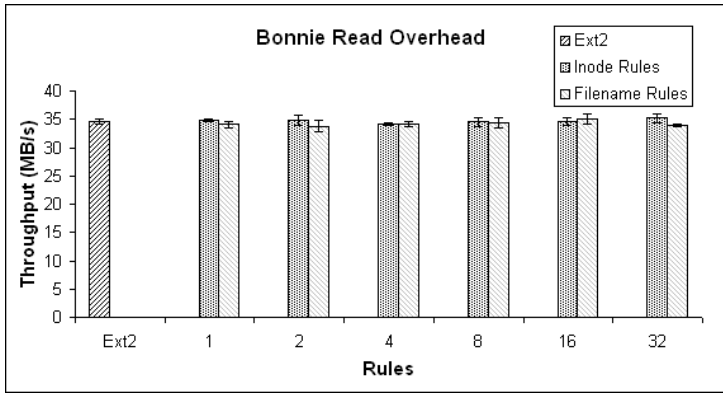


Fig. 5a. Throughput in Bonnie++ benchmark for sequential reads

As can be seen in Figures 5a and 5b, UmbrellaFS imposes no discernable overhead in large sequential reads and writes. The read benchmark in particular was rather noisy, as can be seen from the confidence intervals. The half widths of the intervals in the read test are all less than 10% of the mean, but the means themselves are often within the confidence intervals of other points in the graph. The confidence intervals for the write overhead are several orders of magnitude smaller, and the means are very close to the base ext2 case. We can infer from this that there is not a large overhead associated with large reads and writes.

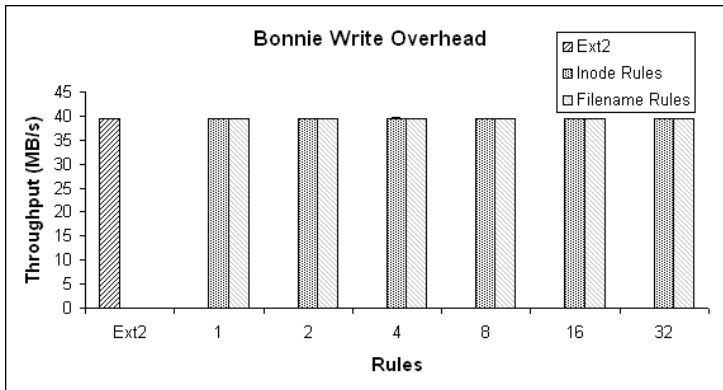


Figure 5b. Throughput in Bonnie++ benchmark for sequential writes

4.2 OpenSSH

Compiling OpenSSH [OpenBSD 2007] is a CPU-intensive benchmark [Wright et al. 2004]. We used a vanilla copy of the OpenSSH 4.7p1 source code. The benchmark itself consists of a typical compilation and installation of OpenSSH, first by running configuration tests, compiling the OpenSSH 4.7p1 code, and finally installing OpenSSH. This is the typical series of configure, make, and make install to compile and install a program. This test involves numerous reads, writes, and other system operations, and helps expose UmbrellaFS's potential impact on typical users. The benchmark was timed, and the results are shown in the amount of time spent in user mode, kernel mode, and wait time, which typically corresponds to I/O, although wait time could also be impacted by many other factors in the system.

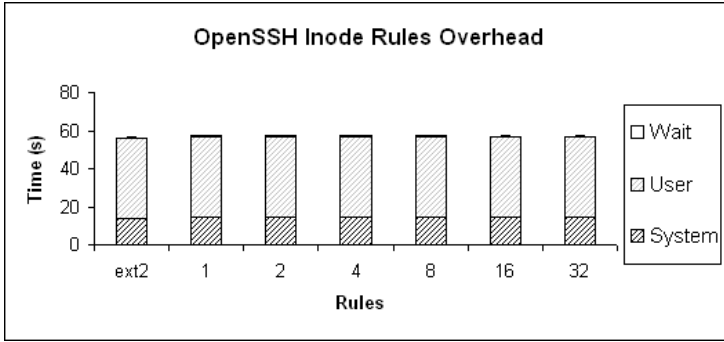


Fig. 6. OpenSSH overhead with inode based rules

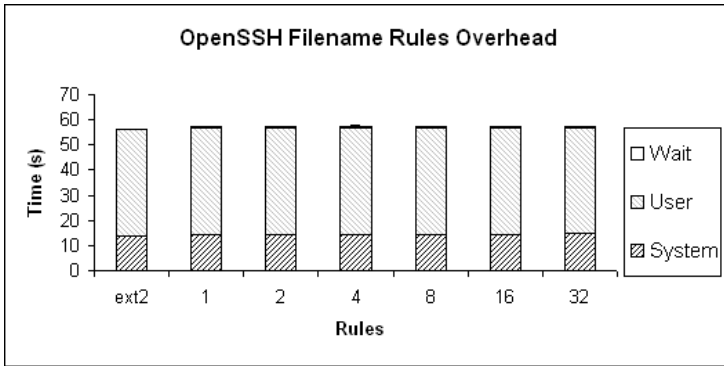


Fig. 7. OpenSSH overhead with filename based rules

Figures 6 and 7 show the overhead of the OpenSSH benchmark for inode and filename based rules. The confidence intervals shown are for overall execution time. In general, the half widths of the confidence intervals were within 1% of the mean, for overall, system, and user times. Wait time had a very large confidence interval, although given the relatively small amount of wait time, this did not have a large effect on the overall execution.

With both inode and file based rules, the predominant overhead is in the kernel execution time, as is to be expected. UmbrellaFS requires an extra system call for every file system operation, since instead of accessing the file systems directly, VFS accesses an UmbrellaFS function. The average kernel execution time in the inode rule test was actually lower for 32 rules than for a single rule, and in the filename rules the overhead was less than 1% higher than the single rule case, indicating good scaling properties for the number of rules.

4.3 Postmark

Postmark is an I/O intensive benchmark designed to simulate the operation of an e-mail server. In our Postmark tests, we used Postmark version 1.5 to create 12,500 files between 8 kB and 64 kB in 200 subdirectories and then performed 25,000 transactions. The block size was 512 bytes, with the default operation ratios and unbuffered I/O.

Figure 8 shows the results of the Postmark tests. Both types of rules showed an overhead on a straight Ext2 partition, typically less than 1% but at no time more than 3%.

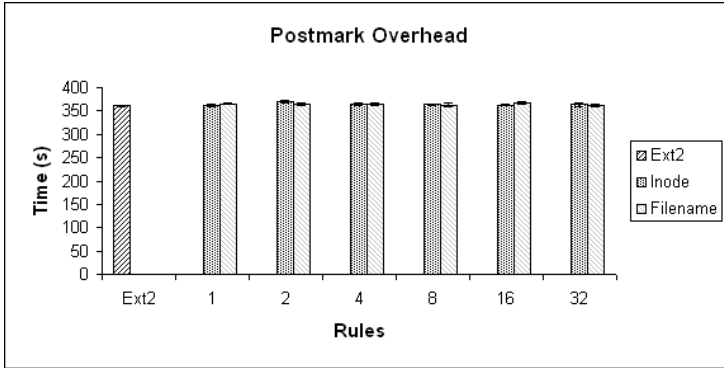


Fig. 8. Postmark Overhead

In the Postmark evaluations, UmbrellaFS showed a consistent overhead over the vanilla Ext2 results. Since Postmark is a benchmark that spends almost all of its time reading and writing to random files, this is not surprising. That both CPU intensive and random I/O intensive benchmarks show less than 3% overhead across the board for UmbrellaFS indicates that in most user applications, the overhead of simply installing UmbrellaFS will not be prohibitively high.

4.4 Overflow

Because the Postmark benchmark had the highest overhead with UmbrellaFS, we also used Postmark to determine the overhead when one file system fills up and another is used in a backup capacity. A single inode-based rule and two branches were used. One file system was filled with data, while the other was empty at the beginning of the test. The rule indicated that writes should go to the full file system first, and then the empty file system. Figure 10 shows the results, compared with vanilla Ext2 and a single inode-based rule on one file system. The overflow system had an execution time very slightly lower than the single rule test, although given the confidence intervals it is not an unreasonable value. The overflow situation showed 2.6% overhead over vanilla Ext2.

As can be seen from Figure 9, simply being forced to write to another drive when the first drive is full imposes no meaningful overhead beyond that imposed by UmbrellaFS in general. In this situation, there was no need to copy files across file systems, since the primary file system began the test full. The overhead of being forced to copy files is evaluated later.

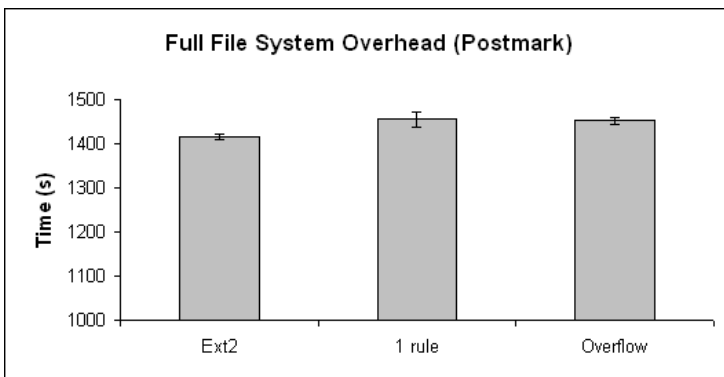


Fig. 9. Postmark overhead with full file system

4.5 Rewriting Overhead

Situations may arise where a file is moved back and forth between two branches on separate drives. While careful application of rules can reduce the risk of this thrashing, it is still conceivable that the properties of a file will hover near the boundary between two rules, requiring multiple rewrites of a file as the characteristics change. In order to examine this situation, we used a micro benchmark which took a set of 35,000 files, appended to each file, and then truncated the file by the same amount, and finally appended again to the file in order to keep it at the larger size. This was repeated 3 times, with the middle set of append, truncate, append taking the file across a rule boundary. The system was set so that files larger than 8 kB were on one file system, and those smaller were on a different file system. Initially files were 6.5 kB in size, and all appends and truncates were 1 kB in size. The results of appends and truncates which moved the file size between 7.5 kB and 8.5 kB are shown in Figure 10. The appends and truncates when not on a rule boundary resulted in overhead in line with that experienced in previous benchmarks, and those results are not shown.

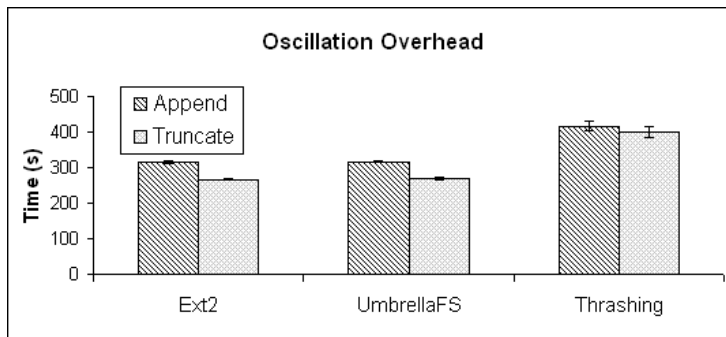


Fig. 10. File size oscillation between 7.5 and 8.5 kB

When a 7.5 kB file was appended with 1 kB of additional data and there was a rule indicating that files larger than 8 kB should be stored on one drive and smaller than 8 kB should be stored on another drive, significant overhead was seen. This is expected, since in a typical append of this size, only an append is done. To move the file, a new file must be created on the other drive, and the full 8.5 kB must be written, rather than just the 1 kB append. The original file must also be deleted. Likewise, on a truncate the file is moved back to the original file system. Appends showed an overhead of 32% on ext2 and 31% over UmbrellaFS without the file movement. Truncates showed an overhead of 49% and 48% respectively. The severity of these overheads demonstrates the importance of careful rule selection. It is noted that rules based on thresholds could be susceptible to this overhead if the workload causes the files to go back and forth across the thresholds, and we are looking to include hysteresis to reduce this type of thrashing.

5. EXAMPLE UMBRELLAFS SCENARIOS

In this section, we demonstrate the potential utility of UmbrellaFS in a number of example scenarios. The first example consists of a system with diverse devices. The second example demonstrates a system where different native file systems are employed below UmbrellaFS. The third example considers different failure semantics and shows the possibility of a Coda [Kistler and Satyanarayanan 1992] like disconnected operation.

5.1 Diverse Devices

In the first example scenario, we consider a system consisting of a Flash drive and a RAID array. Samsung's 32 GB Flash drive (model MCBOE32G8APR-0XA) and 5 of the Seagate Cheetah 10k.7 drives organized as a 4+P software RAID5 array were employed. The characteristics of these two devices are shown below in Figures 11a and 11b.

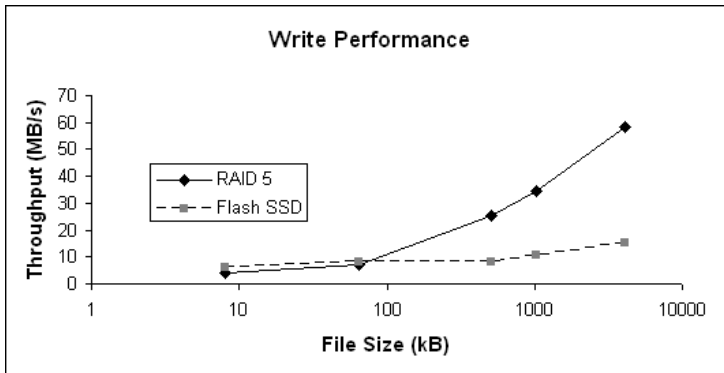


Fig. 11a. RAID 5 and Flash SSD write performance

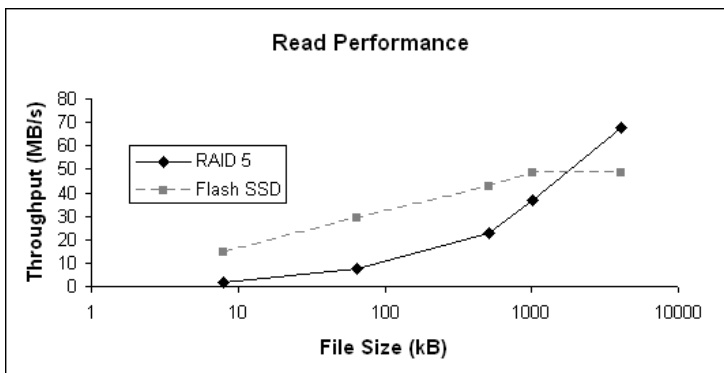


Fig. 11b. RAID 5 and Flash SSD read performance

The Flash drive has superior read performance at request sizes up to 2 MB and better write performance at smaller write sizes up to 64 kB. The RAID array has higher read and write performance at larger request sizes. It is noted that the performance gap of writes between the flash drive and the RAID arrays is significant (up to a factor of 4) at larger file sizes.

In order to mitigate the impact of limited number of write cycles of flash drives and the write performance gap at large file sizes, we decided to allocate read-only files and executable files on the flash drive while placing the remaining files on the RAID array. Postmark was modified to emulate this situation. Two file sets were created. One file set was only read, while the other file set was both read and written to. The policy made this decision based on filenames in this example; however in a real system the file permissions could just as easily be used. File sizes in this simulation were between 8 kB and 1 MB. The performance of such a policy on the example system is shown in Figure 12.

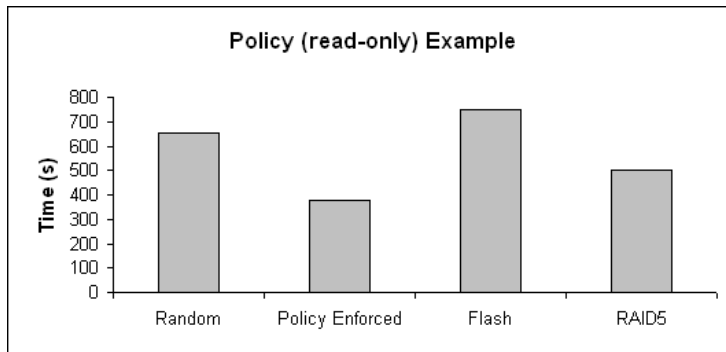


Fig. 12. Diverse Devices Policy Rule Example

The results of this policy are compared against the policy of randomly assigning files across the two types of devices, as well as placing all the files on either the flash drive or the RAID array. The policy took 44% less time than a random distribution, and 51% and 26% less time than flash and RAID 5, respectively. The results show that by keeping writes on the RAID 5 array and heavily utilizing the Flash device for reads, the overall performance is improved. In the random distribution example, the large writes to the slower flash device impacted performance. In this example, the flash drive did not fill up, which could have forced overflow to the RAID array or vice-versa. Even had this occurred, however, because there would not be a need to move the files, the system would not exhibit the overhead of section 4.5, but rather the performance would have reverted to the RAID 5 curves in Figures 12a and 12b. Performance improvements are dependent on many factors, viz., the relative performance of the devices and the relative loads on the different devices, localities of loads and other factors and that performance gains from policy decisions may not be universal [Shenoy et al. 1998].

More importantly, the policy of allocating read-only and executable files to flash drive through a system level policy improves the lifetime of the flash drive since writes are incurred only once when the files are initially written to the flash drive. Second, if the user or administrator deems that small files can also be effectively placed on flash drives (because of RAID array's small-write cost), it is easy to modify or augment the system policy to make this change.

5.2 Diverse File Systems

Another example scenario was set up similar to the previous situation. In this case, the read only file set consisted of 50,000 small files, with sizes ranging from 1 kB to 5 kB, rather than the large files represented in the previous example. The read/write file set still consisted of files between 8 kB and 1 MB. The performance of this situation was examined in terms of the time to create the files as well as the time to run transactions on them (reads in the case of the 50,000 small files, reads, appends, creations, and deletions for the large files). The small files were placed on the Flash drive as before, but this time the file system on the Flash drive was Ext2 for one set of runs and Nilfs2 [Yoshiji et al. 2008], a Linux log-structured file system, for another. The other set of runs simply placed all the files on the RAID 5 system. The results of the test are shown in Figure 13.

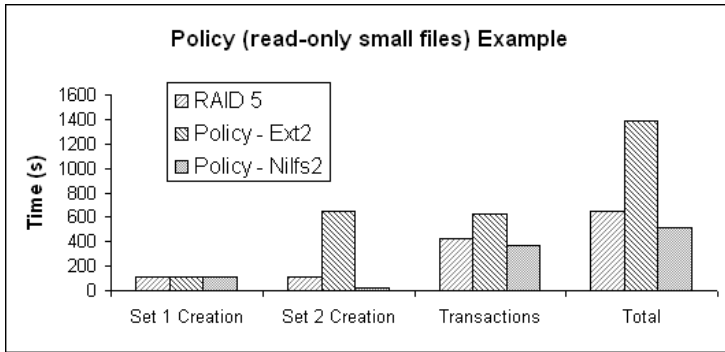


Fig. 13. Diverse Files Systems Policy Rule Example

As can be seen from the results, the Flash device performs poorly with respect to file creation with numerous small files when it uses the Ext2 file system. When using Nilfs2, however, the Flash device took 79% less time to create the files than the RAID 5 array, and over 96% less time than the same Flash device using Ext2. Transactions took 11% less time with the policy based storage and Nilfs2 on the Flash device and 40% less time than the same policy-based storage with Ext2 on the Flash device. These results demonstrate that different file systems can provide a drastic difference to results for different devices. While a traditional file system might work well for a traditional magnetic drive, a log based file system might provide better results for a Flash based device. As this example shows, UmbrellaFS makes matching particular file systems to particular devices possible within a single namespace.

5.3 Encryption

In order to examine a situation with multiple types of file systems, consider the example of a system with 3 types of storage devices, a magnetic disk, a magnetic disk with encryption, and a flash drive.

Difficulties were encountered obtaining Linux drivers for full disk encryption drives such as the Seagate Momentus 5400 FDE.2. So rather than relying on a hardware encryption solution, EncFS [Gough 2006], an open source encrypted file system that makes use of the FUSE [Szeredi 2007] library to run in user space, was used in a software-based emulation. The blowfish algorithm was used to encrypt files with a 160 bit key. While this does not function exactly as a hardware-based encrypting drive would, the overhead from encryption is present, as well demonstrating the functionality of UmbrellaFS working across multiple types of underlying file systems.

The same modified Postmark benchmark as the previous example was used in this situation. Three drives were used, two of which were the Seagate Cheetah drives, as well as the Samsung Flash drive. The numbers in this example and the previous one are not intended to be directly compared.

Figure 14 shows the results from a policy where only the 10% files placed on the “encrypted” drive were encrypted compared to a policy where the files are distributed the same way between the different drives, but all three file systems were encrypted. As is to be expected, when only a portion of the files encounter the overhead of encryption and decryption, the benchmark takes much less time to complete than when all files need to be encrypted before being stored and decrypted before being read. This example demonstrates how appropriate use of rules can reduce the burden of particular file placement on underlying drives on the user. It also shows the ability of UmbrellaFS to function across multiple different underlying file systems.

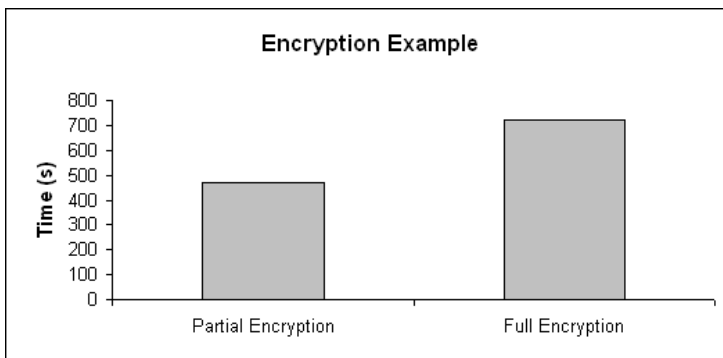


Fig. 14. Encryption Example

5.4 Disconnected Operation

CMU's Coda system achieves disconnected operation in a networked storage situation by caching recently used files on the local machine. When the machine is disconnected, operation can continue on those local files, and then when reconnected, the copies on the network are updated to match the locally cached copies. With appropriate use of rules, UmbrellaFS can achieve a somewhat similar type of operation.

In this situation, assume that the user has a local machine and has mounted an NFS partition as well. A drive on the local machine and the NFS partition can be used to achieve disconnected operation. By setting up a rule based on access time, all files which are accessed today can be moved from the NFS partition to the local machine.

If the NFS server should become unavailable, either due to a network outage, or simply because the user's machine has been disconnected from the network (in the case of say, a laptop that the user is taking home for the day), then operations that attempt to access the directory that UmbrellaFS is maintaining will return an error. If rather than an OR based evaluation of errors (where any error in any of the underlying branches returns up an error to the user), we utilize an AND based evaluation, we can hide this error from the user and take appropriate action to return the available information. The choice of between OR and AND based evaluation of errors can be made when the module is loaded, the same way the rules themselves can be specified.

When an error of this type is encountered, we rely on a feature of Unionfs, that is, the ability to add and remove underlying branches on the fly. By calling an appropriate `ioctl`, we can remove the branch that returned the error from UmbrellaFS, so that future actions will not attempt to access that branch. This functions in the same manner as in Unionfs, although in addition to the work that is normally done, the rules must be updated. If not, then an access that would normally go to the (now disconnected) device would cause problems such as accessing indices that are not present, etc. So when a branch is removed, we analyze the rules in the system, and if any rule's primary file system that it directs to should become unavailable, that rule is flagged as inactive. Inactive rules are then simply skipped by the rule evaluation mechanism. When evaluating the placement of a file, if no rule applies to a file because of a rule deactivation, then just as if there is no rule in the system that applies to that file, the file is placed on a default file system.

When the NFS partition is reconnected, we again utilize the Unionfs ability to add branches on the fly. Again we must evaluate the rules, and if an inactive rule's primary branch has returned, we reactivate that rule so that it again directs files to the newly returned drive. In this situation, however, we are faced with the files that have been moved to the local machine and need to be migrated back to the network storage. In

order to accomplish this, we utilize a background program. Since our evaluation of a file and its correctness of location is done on open and close, we cannot move files back to the networked storage without assistance. The background program simply opens and closes files, so that the evaluation can take place. Upon closing, the most recent access time is evaluated, and if the file has not been accessed recently, then it is moved back to the network storage device. Since the program does not actually read the file, the file's access time is unaffected and can be properly examined.

By using these elements of the UmbrellaFS system, we can achieve disconnected operation similar to that provided by Coda, without requiring a fully new file system. The examples presented in this paper are not intended to be comprehensive, but rather an indication of the types of things that can be achieved with an UmbrellaFS type system.

6. RELATED WORK

IBM's General Parallel File System (GPFS) is a high-performance parallel storage solution. GPFS provides users flexibility in striping policies for individual files. UmbrellaFS can provide some of the same benefits without necessitating a total switch of file systems. UmbrellaFS can direct files to appropriate file systems mapped to the underlying device characteristics.

Duke's Active Names [Vahdat et al. 1999] framework is another system that can be leveraged in a way similar to UmbrellaFS. In this case, however, to be used as a file system, the Active Names system would need to operate as its own file system. In addition, the Active Names framework focuses exclusively on the namespace. UmbrellaFS can make decisions based on both the namespace and other metadata such as file modification times, permissions, and the like.

HP's AutoRAID system is a device level approach to managing storage across RAID 1 and RAID 5 devices. In AutoRAID "hot" and "cold" data are stored in the RAID 1 and RAID 5 arrays, respectively. RAIF [Joukov et al. 2007] enables policy driven striping of files across file systems. Packet interposing [Anderson et al. 2000] has been shown to provide benefits in scaling performance in a clustered file server.

Intel's Robson technology [Trainor 2006] has proposed using NAND Flash RAM in order to operate as a cache for traditional hard drive accesses. By storing files in Flash RAM, the amount of time the system spends rotating disk arms on traditional hard drives can be dramatically reduced.

OSD framework potentially allows storage systems to be tailored to specific characteristics of objects. However, the current OSD architecture or framework leaves the "policy" of allocation open and we are not aware of any specific studies investigating the utility of different allocation policies in OSD systems. Panasas's storage allocation policy scales the striping granularity based on the file size [Panasas 2005]. Policies have been employed in managing grid resources [Bharathi et al. 2005] and in enforcing security [Blaich et al. 2007].

Brick based storage systems such as Self-* [Ganger et al. 2003] and FAB [Frølund et al. 2004] typically distribute file blocks randomly or based on a policy across the available storage bricks or nodes. However, this policy is typically not tailored to nature of the file or other characteristics of the file. The Google File System [Ghemawat et al. 2003] randomly distributes the blocks of a file across a number of nodes.

Sun's ZFS file system [Sun 2007] allocates blocks flexibly across a pool of storage devices at the time of writing the file, rather than tying the file system to a device at the time of creating the file system. However, ZFS doesn't provide the user the flexibility of allocating a file across available devices.

Conquest File System [Wang et al. 2006] is designed to span both persistent memory and magnetic disks. UmbrellaFS can provide similar functionality in terms of the placement of files based on size, but cannot easily place metadata on one drive and data

on another due to metadata consistency and underlying file system issues. UmbrellaFS has numerous additional rule possibilities in addition to Conquest's file-size threshold approach.

The Veritas File System [Symantec 2007] has policies which can direct files to particular underlying storage devices. Unlike UmbrellaFS, the enforcement of these policies to move files between devices is not done on the fly as situations change, but must be directly triggered by an administrator.

7. CONCLUSIONS

In this paper, we presented the rationale and design of Umbrella File System. UmbrellaFS is designed to provide flexibility in exploiting the device level differences that are becoming more prevalent with the advent of Flash based solid-state drives. Our prototype implementation on Linux 2.6 showed that UmbrellaFS adds little overhead to most of the file system operations. We also demonstrated the utility of UmbrellaFS in both an example system organization consisting of a Flash drive and a RAID array, as well as a system utilizing selective encryption of files. We have shown that with appropriate rules, UmbrellaFS can manage storage across diverse devices efficiently, providing performance improvements of up to 44% in a system with both Flash based storage and a RAID array of SCSI disks.

We are looking at providing hysteresis to the system in order to minimize the negative effects of oscillation. We are also looking at the ability to include file system based rules, such as rules concerning the current available space in file systems or other aspects that are not necessarily present in the file metadata. We also plan to study the potential of automatically inferring beneficial policy rules from file access behavior in guiding the policy decisions.

REFERENCES

- ANDERSON, D.C., CHASE, J.S., and VAHDAT, A.M. 2000. Interposed request routing for scalable network storage. In *Proceedings of the Fourth Symposium on Operating Systems Design and Implementation (OSDI)*, San Diego, CA, 259-272.
- BHARATHI, S., KIM, B.K., CHERVENAK, A., and SCHULER, R. 2005. Combining virtual organization and local policies for automated configuration of grid services. In *Proceedings of Self-Organization and Autonomous Systems in Computing and Communications (SOAS) 2005*, Glasgow, UK, 194-202.
- BLAICH, A., LIAO, Q., ALLAN, G., STRIEGEL, A and THAIN, D. 2007. Lockdown: distributed policy analysis and enforcement within enterprise network. In *Proceedings of the 16th Usenix Security Symposium*, Boston, MA.
- COKER, R. 2001. Bonnie++. <http://www.coker.com.au/bonnie++/>.
- FRÖLUND, S., MERCHANT, A., SAITO, Y., SPENCE, S., and VEITCH, A. 2004. FAB: building distributed enterprise disk arrays from commodity components. In *Proceedings of the 11th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, Boston, MA, 48-58.
- GANGER, G.R., STRUNK, J.D., and KLOSTERMAN, A.J. 2003. Self-* Storage: brick-based storage with automated administration. *Carnegie Mellon University Technical Report CMU-CS-03-178*.
- GHEMAWAT, S., GOBIOFF, H., and LEUNG, S. 2003. The Google File System. In *Proceedings of the Nineteenth ACM Symposium on Operating Systems Principles (SOSP)*, Bolton Landing, NY, 29-43.
- HALCROW, M. 2005. eCryptfs: An enterprise-class cryptographic filesystem for Linux. In *Proceedings of the Linux Symposium*, Ottawa, Ontario, Canada, 1:201-218.
- HOHMANN, C. 2007. CryptoFS. <http://reboot.animeirc.de/cryptofs/>.
- JOUKOV, N., KRISHNAKUMAR, A.M., PATTI, C., RAI, A., SATNUR, S., TRAEGER, A., and ZADOK, E. 2007. RAIF: Redundant Array of Independent Filesystems. In *Proceedings of the 24th IEEE Symposium on Massive Storage Systems and Technologies*, 199-214.
- KATCHER, J. 1997. Postmark: A new file system benchmark. Tech. Rep. TR3022, Network Appliance. http://www.netapp.com/tech_library/3022.html
- KISTLER, J., and SATYANARAYANAN, M. 1992. Disconnected operation in the Coda file system. *ACM Transactions on Computer Systems*, 10, 1 (February 1992), 3-25.
- KLEIMAN, S.R. 1986. Vnodes: An architecture for multiple file system types in Sun UNIX. In *USENIX Association: Summer Conference Proceedings*, Atlanta, GA, 238-247.

- MATHUR, G., DESNOYERS, P., GANESAN, D., and SHENOY, P. 2006. Capsule: An energy-optimized object storage system for memory-constrained sensor devices. In *Proceedings of the Fourth ACM Conference on Embedded Networked Sensor Systems*, Boulder, CO, 195-208.
- GOUGH, V. 2006. EncFS Encrypted Filesystem. <http://arg0.net/wiki/encfs>.
- OPENBSD. 2007. OpenSSH. <http://www.openssh.org>.
- PANASAS, INC. 2005. Object Storage Architecture White Paper. http://www.panasas.com/docs/Object_Storage_Architecture_WP.pdf
- PATERSON, D.A., GIBSON, G., and KATZ, R.H. 1988. A case for redundant arrays of inexpensive disks (RAID). In *Proceedings of the 1988 ACM SIGMOD International Conference on Management of Data*, Chicago, IL, 109-116.
- PATERSON, R.H., GIBSON, G.A., and SATYANARAYANAN, M. 1992. Using transparent informed prefetching (TIP) to reduce file read latency. In *Proceedings of the NASA Goddard Conference on Mass Storage Systems*, Greenbelt, MD, 329-342.
- SAMSUNG SEMICONDUCTOR EUROPE. 2007. Samsung Flash Solid-State Drive. http://www.samsung.com/eu/Products/Semiconductor/downloads/Samsung_SSD_for_mobile.pdf
- SEAGATE TECHNOLOGY LLC. 2007. Seagate unveils new giants -- 250GB notebook hard drive and the first encrypting 1TB desktop PC drive. <http://www.seagate.com>.
- SHENOY, P., GOYAL, P., and VIN, H. 1998. Architectural considerations for next generation file systems. *University of Massachusetts Technical Report UM-CS-1998-048*.
- SUN MICROSYSTEMS. 2004. ZFS: the last word in file systems. <http://www.sun.com/2004-0914/feature/index.html>.
- SZEREDI, M. 2007. Filesystem in Userspace. <http://fuse.sourceforge.net/>.
- SYMANTEC. 2007. HP and Symantec: Enabling ILM solutions with dynamic storage tiering. Symantec White Paper, <http://eval.symantec.com>.
- TRAEGER, A., ZADOK, E., JOUKOV, N., and WRIGHT, C.P. 2008. A nine year study of file system and storage benchmarking. In *ACM Transactions on Storage (TOS)*, 4, 2 (May 2008), 1-56.
- TRAINOR, M. 2006. Overcoming disk drive access bottlenecks with Intel® Robson technology. In *Technology@Intel Magazine*, 4, 9, 9-11.
- VAHDAT, A., DAHLIN, M., ANDERSON, T., and AGGARWAL, A. 1999. Active Names: flexible location and transport of wide-area resources. In *Proceedings of the USENIX Symposium on Internet Technologies and Systems (USITS)*,
- WILKES, J., GOLDING, R., STAELIN, C., and SULLIVAN, T. 1996. The HP AutoRAID hierarchical storage system. In *ACM Transactions on Computer Systems* 14, 1, 108-136.
- PROJECT T10. 2004. Information technology—SCSI object-based storage commands (OSD), Rev. 10.
- WANG, A., KUENNING, G., REIHER, P., and POPEK, G. 2006. The Conquest file system: better performance through a disk/persistent-RAM hybrid design. In *ACM Transactions on Storage (TOS)*, 2, 3, 309-348.
- WRIGHT, C.P., DAVE, J., GUPTA, P., KRISHNAN, H., ZADOK, E., and ZUBAIR, M.H. 2004. Versatility and Unix semantics in a fan-out unification file system. Stonybrook University Technical Report FSL-04-01b.
- YOSHII, A., KONISHI, R., SATO, K., HIFUMI, H., TAMURA, Y., KIHARA, S., MORIAI, S. 2008. New Implementation of a Log-structured File System, version 2.0.6. <http://www.nilfs.org/en/index.html>
- ZADOK, E., IYER, R., JOUKOV, N., SIVATHANU, G., and WRIGHT, C.P. 2006. On incremental file system development. In *ACM Transactions on Storage (TOS)*, 2, 2, 1-33.