SiRiUS: Securing Remote Untrusted Storage

Eu-Jin Goh[*], Hovav Shacham[†], Nagendra Modadugu, Dan Boneh[‡]
Stanford University
{eujin, hovav, nagendra, dabo}@cs.stanford.edu

Abstract

This paper presents SiRiUS, a secure file system designed to be layered over insecure network and P2P file systems such as NFS, CIFS, OceanStore, and Yahoo! Briefcase. SiRiUS assumes the network storage is untrusted and provides its own read-write cryptographic access control for file level sharing. Key management and revocation is simple with minimal out-of-band communication. File system freshness guarantees are supported by SiRiUS using hash tree constructions. SiRiUS contains a novel method of performing file random access in a cryptographic file system without the use of a block server. Extensions to SiRiUS include large scale group sharing using the NNL key revocation construction. Our implementation of SiRiUS performs well relative to the underlying file system despite using cryptographic operations.

1. Introduction

Remote file storage is common in many different environments: in large organizations, using protocols such as NFS [32, 5] and CIFS [22]; in P2P networks, using systems such as OceanStore [27], Farsite [23], and Ivy [25]; on the web, using services such as Yahoo! Briefcase [34]. In all these environments, the end-user often has no control over the remote file system. Many of these storage systems rely on the untrusted remote server for data integrity and access control. However, the remote server often does not guarantee data integrity and only implements weak access control.

In this paper, we design and implement a security mechanism that improves the security of a networked file system *without making any changes to the file system or network server*. Our motivation for insisting on no changes to the underlying file system is twofold.

First, we want a system that is easy for end-users to install and use. Since end-users often have no control of the remote server, they cannot install a security system that requires changing the remote file server. Similarly, organizations who have invested in large NAS devices have no control over the inner implementation of the device. These organizations can only add security mechanisms that do not require changing the NAS server. Hence, insisting on no changes to the file server enables us to enhance the security of legacy network file systems without changing the existing infrastructure.

Second, we want a security mechanism that can be layered on top of any network file system; NFS, CIFS, P2P, HTTP, etc. By insisting on no changes to the underlying file systems, SiRiUS becomes agnostic of the network storage system. The result is a uniform mechanism for securing many different types of network storage.

Our insistence on making no changes to the remote server comes at a cost. We cannot defend against certain attacks such as denial of service — an attacker capable of compromising the remote server can delete files. However, the attacker will not be able to view or undetectably alter files. Other security implications of this design principle are discussed in Section 5. For this reason we view the SiRiUS system as a stop-gap measure designed to add security to legacy systems. The importance of securing legacy systems stems from people's reluctance to upgrade. This is nicely captured in the following quote from the designers of NFSv3 [4, p. 61] four years after NFSv3 was introduced:

> Although NFS version 2 has been superseded in recent years by NFS version 3, system administrators are slow to upgrade the operating systems of their clients and servers, so NFS version 2 is not only widespread, it is still by far the most popular version of NFS.

SiRiUS is designed to be easy to install and is intended to be used until a new network file system is deployed with adequate access control and data integrity abilities.

SiRiUS is designed to handle multi-user file systems where users frequently share files. It supports granting read only or read-write access to files. This flexibility is unusual in cryptographic file systems (e.g., CFS [2])

[*] Supported by NSF Grant No. CCR-0205733.
[†] Supported by NSF Grant No. 0121481.
[‡] Supported by NSF Grant No. CCR-0205733.

where possession of a key usually enables both reading and writing to a file. SiRiUS also defends against version rollback attacks as described in Section 3.3 and 6.4. We reduce network traffic by providing random access within files while still ensuring data integrity. SiRiUS is built using the SFS toolkit [17].

The next section describes the design goals of the SiRiUS system. Sections 3 and 4 describe the inner workings of the system and implementation details. Our experiments with SiRiUS show reasonable performance. For read operations, we get a slowdown factor of 2 when using SiRiUS layered over NFS, versus plain NFS. Detailed experimental results are given in Section 4.8. We discuss related work in Section 7. We note that the issue of building an easy-to-deploy secure file system that interoperates with legacy infrastructure seems to have received little attention so far.

2. Design Criteria

In this section, we describe the criteria used in designing SiRiUS. We begin with the general systems requirements and proceed to the security requirements.

2.1. System Considerations

No changes to File Server. SiRiUS must add security to existing network file systems with no change to the software or hardware of the servers. This consideration is fulfilled by layering SiRiUS over existing network file systems. Existing secure file systems (such as SUNDR [19]) that work at the file system block level require extensive modification to the file server. Therefore, SiRiUS should work at a higher level, using files as its atomic unit.

Fulfilling this requirement makes it possible for users to install SiRiUS without the support of the file server administrator. Similarly, it makes it possible to use SiRiUS to secure different types of network storage systems (NFS, CIFS, P2P, Web). In fact, SiRiUS can also be used to secure removable storage devices such as USB hard drives and compact flash devices. Requiring zero changes to the file server limits the security we can provide, as discussed in Section 5.1. For example, we cannot prevent a "sledgehammer" denial-of-service attack, in which an administrator deletes all files.

File Sharing. The ability to share files amongst users is essential in a network file system. SiRiUS users must be able to share a file easily with other users of the system. Existing cryptographic file systems [2, 35, 1, 7, 13] limit their own usefulness because they either provide very coarse sharing at the directory or file system level or fail to distinguish between read and write access. File systems that do provide per-file sharing [24, 19, 29, 18] rely on a trusted authentication mechanism residing on the file server, which precludes their use in settings where users have no administrative control over the server.

Minimal Client Software. A SiRiUS user should need to run only a user-level daemon. Users should not be required to upgrade or patch the client OS kernel.

Performance. SiRiUS should not perform unreasonably worse than its underlying file system.

Efficient Random Access and Low Bandwidth. Reads and writes to any location in a file should take comparable amounts of time. Efficient random access is hard to achieve in a cryptographic file system that does not operate at the file block level. Random access allows partial file reads or writes which take up less network bandwidth.

2.2. Security Considerations

Confidentiality and Integrity. File data must be protected from users that are not granted access. Even the super user administering the remote file server should not be able to read files unless explicitly given permission by the file owner. Furthermore, unauthorized modifications to file data must be detected by SiRiUS.

It is desirable to protect filenames since they may leak information about the contents of a file. However, other pieces of meta data information such as data block pointers and modification times should not be protected to provide better crash recovery. For example, if data block pointers are encrypted, the file system cannot piece together a partially written file.

Untrusted File Server. A serious problem with legacy network file systems is that their access control mechanisms are insecure and easily defeated. For example, the most popular authentication mechanism used by NFS trusts the user and group ID of incoming file requests. These IDs are easily spoofed by an attacker to gain full access to a user's files. Hence, SiRiUS cannot depend on any access control enforcement by the file server.

SiRiUS must store all access control information (encrypted and signed) together with the file data. Storing all access control data together with the file data has the added benefit of making files readily available even after backups, mirroring and replication of the network file server. Furthermore, it also allows file sharing with minimal out-of-band communication with other users.

All file data is encrypted and signed on the client before being stored on the file server (providing end-to-end security). Performing all cryptographic operations on the client has the benefit of offloading the cryptographic

workload from the file server onto lightly loaded client machines. It also obviates the need for channel security when transmitting files over a network.

File Access Controls. SiRiUS must support two file access modes; read only and read-write. Support for write only mode is preferable but not a requirement as it is uncommon. Access control should be enforced on a per file basis. The access control mechanism in SiRiUS should not limit the flexibility of file sharing provided by the underlying file system while still providing security. SiRiUS should also be able to enforce access control even when the underlying file system does not support file sharing. This is easily achieved when all users of a file masquerade as one user (from the server's point of view).

Based on usage patterns in our department file servers, we observe that most files are shared only among a small number of users. Thus, it is reasonable to optimize SiRiUS sharing for small groups.

Key Management. The proliferation of keys used by different applications creates a management and usability nightmare. To avoid compounding this problem, a SiRiUS user must not be required to keep more than a few keys for file system access. These keys must also be compatible with simultaneous use in other applications.

Key Distribution and Revocation. Many cryptographic file systems discount the problem of key distribution and revocation by assuming timely and efficient out-of-band communication between users. We feel that such assumptions are unrealistic and we address such problems directly in SiRiUS.

All out-of-band communication for obtaining other users' keys should be minimized. Ideally, SiRiUS can use existing key distribution infrastructure such as PGP public key servers or even Identity Based Encryption (IBE) [3] master key servers.

In SiRiUS, user access revocation must not use online or out-of-band interaction between users. The user access revocation mechanism for the file system should be simple and effective. When read or write access to a file is revoked, the revoked user should immediately lose access to that file without need for communication.

Freshness Guarantees. SiRiUS must guarantee the freshness of meta data containing the access control information. This freshness guarantee allows timely revocation of access controls and also prevents access control rollback attacks. The freshness of file data should also be guaranteed to ensure that users always read the latest version of their files.

3. File System Design

In this section, we describe the SiRiUS file system data structures and algorithms.

3.1. Overview

Existing secure file systems which support file sharing only consider the problem of securing and sharing files on a trusted file server. Furthermore, many of these designs require custom file servers. The goal of SiRiUS is to secure data placed on any untrusted and unmodified network file server while maintaining performance and standard file system semantics. We next give a brief overview of SiRiUS.

From the SiRiUS user's point of view, SiRiUS appears as a local file system with the standard hierarchical view of files and directories. A SiRiUS client on the user's machine intercepts all operations to the SiRiUS file system and processes the requests before transmission to the remote file server. The type of network file system exported by the remote file server is hidden from the user. All cryptographic operations including encryption and signing are done by the client before the results are placed on the file server.

All SiRiUS users maintain one key for asymmetric encryption and another for signatures. We call these the user's master encryption key (*MEK*) and master signing key (*MSK*).

Files stored on the file server are kept in two parts. One part contains the file meta data and the other the file data. We will refer to the meta data file as *md-file* and the encrypted data file as *d-file*. The file meta data contains the access control information while the file data contains the encrypted and signed contents. The file data is encrypted with a symmetric cipher using a unique key for each file. We call this encryption key the file encryption key (*FEK*). The data is also signed using a signature scheme with a unique key for that file. The signing key is called the file signature key (*FSK*).

The *FEK* and *FSK* are used to differentiate between read and write access. Possession of only the *FEK* gives read only access to the file while possession of both the *FEK* and *FSK* allows read and write access. For example, a user with only the *FEK* cannot create a valid file because he cannot produce a valid file signature.

A key distribution mechanism (such as a PKI) may be required for SiRiUS depending on the encryption and signature scheme used. Organizations that are concerned about security tend to have such infrastructure already in place. Identity Based Encryption [3] and Signature schemes tend to require less infrastructure than a traditional PKI and are well suited for individual and small-scale use. SiRiUS can make use of any existing key distri-

Acronym	Definition
FEK	File Encryption Key
FSK	File Signature Key
MEK	Owner's Master Encryption Key
MSK	Owner's Master Signature Key
d-file	Data File
md-file	Meta Data File
mdf-file	Meta Data Freshness File

Table 1. Glossary.

bution mechanisms that already exist for PGP, S/MIME, IBE, etc.

There is a meta data freshness file (*mdf-file*) located in every directory of a user's file system. This file contains the root of a hash tree [21] built from all the *md-files* in the directory and its subdirectories. For example, the *mdf-file* at the root of a user's home directory contains the root of the hash tree constructed from the user's *md-files* in the directory and *mdf-files* under immediate subdirectories. A user's SiRiUS client will periodically time stamp the root *mdf-file* and sign it using his *MSK*. The update interval can be set by the user.

In the subsequent sections, we describe the file formats and how basic file operations — create, read, write and sharing — are carried out. A glossary is provided for reference in Table 1. For clarity, assume that SiRiUS uses RSA for asymmetric encryption, AES for symmetric encryption, SHA-1 for hashing and DSA for signing.

3.2. File Structure

We first describe the structure of the *md-file*. The *md-file* contains access control information and its format is depicted in Figure 1. Each encrypted key block corresponds to a user with some access rights to the *d-file*. Encrypted key blocks contain the *FEK* encrypted under the *MEK* of each user with read access. If a user also has write access, then the *FSK* is included in the user's encrypted key block. Figure 2 shows examples of key blocks for two users, one with read-write access and the other with only read access. Each encrypted key block is tagged with the user name or key ID corresponding to the public key used to encrypt the block.

The *md-file* also contains the public portion of the *FSK* in the clear so that readers can verify the integrity of the *d-file*. The timestamp is updated by the owner when the meta data file is modified. The *md-file* also contains the relative filename (as opposed to absolute pathname) of the file.[1] Finally, the *md-file* is signed using the owner's *MSK*.

[1]The filename is included in the *md-file* to prevent file swapping attacks. We provide more detail in Section 3.3.

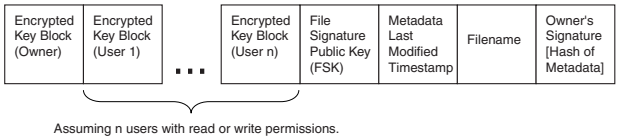

Figure 1. Meta Data file format.

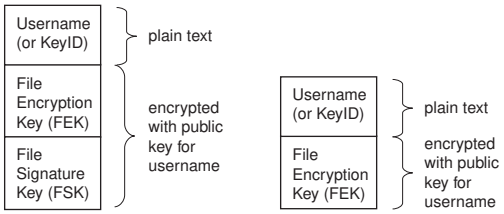

Figure 2. Encrypted Key Block format.

Note that the *md-file* is signed by the file owner's *MSK* and hence can be updated only by the owner. Also note that the first encrypted key block is always encrypted under the file owner's *MEK*.

The *d-file* contains the file data and is shown in Figure 3. File data is encrypted using the *FEK* contained in the corresponding *md-file*. A hash of the data is computed and signed using the *FSK* also contained in the *md-file*. This signature is appended to the end of the file.

3.3. Freshness Guarantees

SiRiUS enables a user to guarantee the freshness of the *md-files* that he owns. Freshness guarantees are required in order to prevent rollback attacks. A rollback attack involves misleading users into accessing stale data. For example, suppose Bob revokes Alice's permission to write to a file named `foo`. Alice does a rollback attack by replacing the new *md-file* with an older version that she saved. The older version of the *md-file* has a valid signature and will hence verify correctly. Alice has now successfully restored her own write permissions to the file. Checking the meta data for freshness would stop such an attack.

SiRiUS uses a hash tree [21] to guarantee freshness. The SiRiUS client for a user generates a hash tree consisting of all his *md-files*. Every directory contains a file with the hash of the *md-files* in that directory and its subdirectories. This file is known as the directory meta data freshness file or *mdf-file*. The directory *mdf-file* is an op-

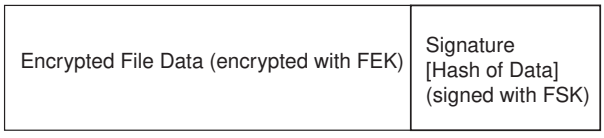

Figure 3. Data file format.

timization for reducing the cost of updating and verifying the hash tree. A similar system for tamper detection in untrusted databases is described by Maheshwari et al. [16].

Generating the *mdf-file*. We next describe how a user's SiRiUS client generates the hash tree for his files and directories. Without loss of generality, we assume that all files and sub-directories (and their contents) contained in a user's home directory belong to the owner. We first describe how the hash tree is generated for a directory.

1. Apply a SHA-1 hash to each meta data file in any order and keep track of the hashes. Also verify the signature on each meta data file during this process.
2. Concatenate the hashes of each meta data file, together with the *mdf-files* of each sub-directory (if any) in lexicographical order and apply a SHA-1 hash to the concatenation.
3. The final hash and the directory name is placed in the *mdf-file*.

Generating the *mdf-file* for a user's root directory is slightly different from the procedure described above. The difference is that the current timestamp is appended to the concatenation of the hashes in the second step when the final hash is calculated. The final hash is concatenated with the timestamp and placed in the root *mdf-file*. The root *mdf-file* is signed with the user's *MSK*.

Verifying Meta Data Freshness. We next describe how Bob's SiRiUS client uses the root *mdf-file* to check the meta data freshness of a file named `foo` owned by Alice.

1. Regenerate the *mdf-file* for the directory where `foo` resides (follow the first two steps of the generation procedure). Compare the result with the current *mdf-file*. If it does not match, then verification has failed.
2. If `foo` is in the root directory, check the *mdf-file* timestamp to ensure it has been updated recently and verify the owner's signature.
   Otherwise, recursively walk up the directory tree and at each step, generate the *mdf-file* and carry out Step 1 till the root directory is reached.

The verification process guarantees that the meta data files in the current directory are fresh up to the timestamp on the root *mdf-file*. Note that it is sufficient to just regenerate hashes for the subtree containing `foo`.

The directory *mdf-files* greatly reduces the cost of regenerating the hash tree because it removes the need to descend into subdirectories that are not in the path to the root. The verification procedure can be accelerated by caching the hashes of directories contents when they are first obtained. The cached values can be used as long as contents of *md-files* remain unchanged.

Updating the Root *mdf-file*. The root *mdf-file* is periodically updated with a new timestamp and signed by the owner using a freshness daemon. For efficiency, the owner should keep a local cached copy of the current root *mdf-file*. Since only one signature is required to update the root *mdf-file*, the update operation is cheap. Hence, the update time interval can be on the order of minutes or seconds.

If any meta data file is changed by the owner, the hash tree needs to be updated. Note that changes to a meta data file only affect the hash calculations of the directory containing the file and the parents of that directory. Hence, we only need to recompute the hash tree for those directories. The same verification procedure is used to regenerate the root *mdf-file* on an update. The differences are that the existing directory *mdf-files* are replaced with the recalculated versions while traversing up the tree. Also, the root *mdf-file* is updated with the current timestamp and re-signed.

Including Filenames in *md-files*. A certain class of attacks cannot be prevented by checking just the meta data freshness. A *md-file* has to be tightly linked to its file name to prevent a file swapping attack.

We provide an example of a file swapping attack. Suppose user Alice owns file `foo` and file `bar`, which reside in the same directory. User Bob wants to read `bar` but has no read or write access to `bar`. However, Bob has read access to `foo`. Bob can trick Alice into writing to the wrong file in the following manner: Bob renames `bar` to `foo` and vice versa, along with their *md-files*. Bob can now read the file named `bar` (which was originally `foo`) but not the file named `foo` (which was originally `bar`). At some point, Alice writes her nuclear launch codes to what she thinks is `bar` without checking the contents of the file. Bob notices the update, reads the launch codes and launches some nuclear missiles.

If the filename were not included in the *md-file*, this attack would succeed because Alice cannot easily verify that she is writing to the right file. Observe that the freshness hash tree would still verify. To prevent this attack, we include the filename in the *md-file*. We also include the directory name in the *mdf-file* to prevent directory swapping attacks. Note that we only need to include the relative filename in the *md-file* and *mdf-file* because the position of the file in the freshness hash tree gives sufficient information about the path of the file.

3.4. Creating a File

A file is created by the SiRiUS client for a user in the following steps.

1. Generate a random DSA File Signing Key (*FSK*) and a random AES File Encryption Key (*FEK*).
2. Encrypt the *FSK* and *FEK* using RSA with the

owner's *MEK* and tag the cipher text with the owner's user name to form the encrypted key block.
3. Apply SHA-1 to the encrypted key block, public key of the *FSK*, a timestamp (of the current time), and filename. Sign the hash with DSA using the user's *MSK*.
4. Create the *md-file* by concatenating the owner's encrypted key block, public key of the *FSK*, the timestamp, the filename, and the signature.
5. Encrypt the file data with AES using the *FEK*. Apply SHA-1 to the encrypted file data and sign the hash with DSA using the private key of the *FSK*. Concatenate the cipher text with the signature to create the *d-file*.
6. Update the root *mdf-file*.

3.5. File Sharing

Recall that only the file owner can permit other users to gain access to the file. SiRiUS sets access permissions by adding encrypted key blocks in the *md-file* for the users the owner wishes to give access to. The following procedure is carried out by SiRiUS to share a file owned by Alice for a user with the user name Bob.

1. Alice obtains the *md-file* and verifies the signature with her *MSK*.
2. Alice obtains the public key for Bob through a public key server.[2] If Bob is only granted read access, Alice encrypts only the *FEK* using RSA with Bob's public *MEK*. If Bob is also granted write access, Alice encrypts both the *FEK* and *FSK*. The cipher text, together with Bob's user name is the encrypted key block to be added to the *md-file*.
3. Alice adds Bob's encrypted key block to the *md-file* and updates the timestamp to the current time. She applies SHA-1 to the modified *md-file* and signs the hash using DSA with her *MSK*. She replaces the signature on the *md-file*.
4. Alice simultaneously verifies the freshness of the old *md-file* and calculates the new root *mdf-file*. If freshness is verified, she updates the root *mdf-file* and replaces the old *md-file* with the new version.

3.6. Writing to a File

The SiRiUS client takes the following steps to write to a file.

1. Obtain the *md-file* and identify the file owner by extracting the user name tag from the first encrypted key block. Obtain the owner's *MSK* using a public key server[3] and verify the signature on the *md-file*. Also verify the freshness of the *md-file*.
2. Locate the encrypted key block with the writer's user name in the *md-file* and decrypt the key block to obtain the *FEK* and *FSK*.
3. Obtain the *d-file* and verify the signature using the public key of the *FSK*.
4. Decrypt the encrypted file data with the *FEK*. Carry out the file write on the plain text file data. If the file is to be replaced, then this decryption step is unnecessary.
5. Encrypt the modified file data with the *FEK*. Evaluate the SHA-1 hash of the encrypted file data and sign the hash with the *FSK*. Append the signature to the newly generated cipher text to create the new *d-file*. Replace the old *d-file* with the new version.

Observe that the freshness hash tree is not updated on a file write.

3.7. Reading a File

The SiRiUS client takes the following steps to read a file.

1. This step is identical to Step 1 for file writing.
2. Locate the encrypted key block with the reader's user name in the *md-file* and decrypt the key block to obtain the *FEK*.
3. Obtain the *d-file* and verify the signature using the public key of the *FSK*.
4. Decrypt the encrypted file data with the *FEK* to obtain the file contents.

3.8. Renaming Files and Directories

A file rename requires changing the filename stored in the *md-file* for that file. A directory rename requires changing the directory name stored in the *mdf-file* for that directory. Both types of renames require updating the freshness hash tree. As a result, only the file or directory owner can perform renames. This file system semantic causes some problems with traditional applications that use a rename paradigm to guarantee atomic file updates. We present a solution to this problem in Section 6.4.

On a file rename, if the destination filename exists, then the destination file is overwritten with the source file. The destination file acquires the permissions of the source file. The source file is deleted and the freshness hash tree is updated to reflect these changes.

On a directory rename, the *mdf-file* for that directory is updated with the new directory name followed by a normal freshness hash tree update.

[2] Note that if Boneh-Franklin IBE [3] is used for asymmetric encryption, Bob's public key can be obtained without contacting a key server.

[3] If an Identity Based signature scheme is used, verifying the file owner's signature can be done without a public key server.

3.9. File Links

SiRiUS supports symbolic links if they are also supported by the underlying file system. Symbolic links on Unix file systems are typically implemented as normal files with the file data containing the path of the link. In this case, when the SiRiUS client accesses a symbolic link, it decrypts the contents and obtains the file pointed to by the link.

One file can have two hard links with different names referencing the file. Therefore, SiRiUS cannot support hard links because each *md-file* only contains one filename.

3.10. Key Management

Key management in SiRiUS is very simple because each file keeps track of its own file keys for access control. All users only need to keep track of two keys; the *MSK* and the *MEK*. There is no out-of-band communication if Identity Based encryption and signature schemes are used. Otherwise, a small amount of out-of-band communication is required in order to obtain public keys.

3.11. Key Revocation

Key revocation in SiRiUS is simple and does not require out-of-band communication. For read access revocation, the owner generates a new *FEK*. Using this key, the *d-file* is updated by encrypting the file data with the new key and then signed (using the old *FSK*). The revoked user's encrypted key block is removed from the *md-file* and all the remaining key blocks are updated with the new *FEK*. Finally, the *md-file* is signed with the owner's *MSK* and the root *mdf-file* is updated.

Write access revocation is the same as read access revocation except that a new *FSK* is also generated. The encrypted key blocks are updated with this new *FSK* and the *d-file* is signed with this new key. Note that revoking write access also involves creating a new *FEK* and re-encrypting the data because write access implicitly provides read access.

3.12. Backups

A system administrator can backup the remote file server by using the standard backup tools such as `dump` or `tar`. Note that the administrator does not need to use SiRiUS to perform the backup. Furthermore, the administrator gains no access to the file data. SiRiUS users only need to backup their *MSK* and *MEK*. They can access any backup of their files with just these two keys.

4. Implementation and Performance

In this section, we describe an instantiation of the SiRiUS file system layered over NFS.

4.1. Overview

We implemented a SiRiUS client on Linux that layers SiRiUS over NFS version 3 using the SFS toolkit [17] and OpenSSL [28]. The SFS toolkit provides support for building user-level NFS loopback servers and clients. The SiRiUS client contains a user-level NFS loopback server to communicate with the client machine and a NFS client to communicate with the remote NFS file server. Using a user-level NFS loopback server to interface with the client machine ensures portability because most modern client operating systems have kernel-level NFS clients.

Note that SiRiUS' NFS client can be replaced with a client that supports a different protocol. For example, one can retain the NFS loopback server and replace the NFS client with a Yahoo! Briefcase (HTTP) client. This will allow Yahoo! Briefcase to look like an NFS server to the user.

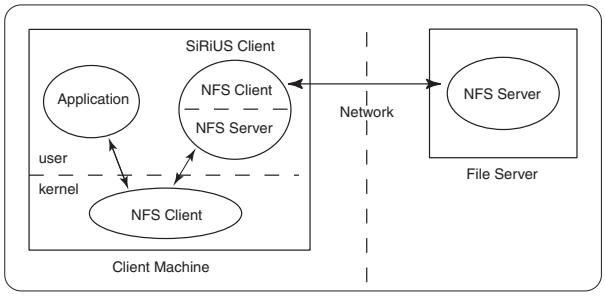

Figure 4. Architecture of SiRiUS layered over NFS using the SFS toolkit.

The SiRiUS client intercepts NFS requests on the NFS-mounted file system and processes the requests. The client then sends the modified requests to the remote NFS server. See Figure 4 for an overview of the system architecture.

4.2. Multiplexing NFS Calls

Many SiRiUS operations require the SiRiUS client to transform a single NFS request from the client machine into multiple requests to the server. First, SiRiUS needs to read and write the *md-file* as well as the *d-file*. Second, many SiRiUS file operations verify the meta data freshness, which necessitates sending a number of NFS calls to walk the directory structure and read *mdf-files*.

The SFS toolkit's ability to perform asynchronous Remote Procedure Calls (RPC) [33] proved a great help in multiplexing NFS calls from the client machine. When the SiRiUS client receives an NFS call from its loopback server, it asynchronously sends out a series of NFS calls to the remote server to process the incoming call. We illustrate this with an example of an NFS CREATE call. When the SiRiUS client receives a CREATE call for a file

named `foo`, it first creates the *md-file*. The SiRiUS client prepends `.x-md-x.` to `foo`, obtaining `.x-md-x.foo` as the name of the *md-file*.[4] It then sends a CREATE call to the remote server for the *md-file*. When this completes, the SiRiUS client generates and stores the contents of the *md-file* using a WRITE call. Following this, the root *mdf-file* is updated. It finally sends the original CREATE request to the remote server and returns the result of that request to the local client machine.

4.3. File System View

The SiRiUS client implementation hides the presence of the *md-files* and *mdf-files* from the SiRiUS user's file system view. For example, an `ls` invocation does not display these meta data files. This view is implemented by processing the NFS READDIR and READDIRPLUS results from the remote server to remove entries for files whose names begin with the meta data file extensions. The code infrastructure added to handle these two NFS calls is easily extended to handle encrypted filenames.

4.4. NFS File Handle Cache

In NFS, all file system objects are identified by a unique NFS file handle generated by the remote server. The SiRiUS client must correlate the handles of *d-files* with those of *md-files*. For example, NFS READ and WRITE operations refer to the target file by its NFS file handle. With just this file handle, we have no way of obtaining the file handle for the associated *md-file*.

Fortunately, the NFS LOOKUP operation is always called on an unknown object to obtain its file handle before other operations can be performed. The NFS LOOKUP call for an object contains the NFS file handle for its directory, and the name of the object in that directory. The LOOKUP call expects a NFS file handle to the specified target in the return result. By caching the arguments and the results of LOOKUP calls, we can maintain enough state to relate *d-file* and *md-file* file handles.

For each file system object, the SiRiUS client caches: its file handle; the file handle of the directory in which it resides; and its name in that directory. This cache is cross-referenced in two hash tables, one keyed by the NFS file handles of objects, and the other by a combination of the directory file handle and directory entry name. The SiRiUS client monitors all NFS operations (e.g., RENAME and REMOVE) that might change NFS file handle state for an object and updates both hash tables.

We give an example to illustrate how file handle caching works. Suppose the SiRiUS client receives an NFS READ call for a file `foo`. It uses the file handle for `foo` given in the READ argument to obtain (from the file handle hash table) the directory handle for the directory in which `foo` resides. Using this directory handle, the SiRiUS client determines the *md-file* file handle (from the directory file handle hash table) using the directory handle and the *md-file* name. The SiRiUS client issues a NFS READ for the *md-file* using this *md-file* handle. In the meantime, the original NFS READ call for `foo` is also sent.

4.5. Changing Access Controls

At present, we have implemented all of SiRiUS described in Section 3. The hooks for adding and revoking permissions to a file are present in the SiRiUS client, but, since SiRiUS permissions are more expressive than Unix permissions, there is no natural way to invoke these hooks directly using the `chmod` system call.

This problem can be solved by a user-level permissions tool that interacts with the SiRiUS client over a dedicated RPC channel. Alternately, the RPC channel can be avoided if the permissions tool calls `chmod` with modified arguments. We illustrate the operation of the permissions tool with an example. Suppose Bob wishes to grant Alice permission to read file `foo`. Bob invokes the command-line tool, which creates a dummy file with a special name. The tool then performs a `chmod` on the file, causing the kernel NFS client to send the NFS SETATTR (set file attributes) request to the SiRiUS client. The dummy filename starts with a special flag and contains the filename `foo` and user name Alice. The SiRiUS client parses the filename, obtains the public key for Alice and performs the appropriate set of NFS calls to change the permissions for file `foo`.

4.6. Random Access and Low Bandwidth

We originally did not plan on implementing random access. While building the SiRiUS client, we realized that whole-file read and write operations provide unacceptable performance for large files. Random access from the SiRiUS client to the remote server is critical when the SiRiUS client must handle partial read and write requests from the local in-kernel client, as in NFS.

The insight is that NFS READ and WRITE calls operate on chunks of 8192 bytes.[5] Hence, reading (or writing) a file larger than 8192 bytes involves multiple NFS operations. If SiRiUS does not have random access, each READ (respectively, WRITE) request involves fetching the entire file to decrypt and verify (respectively, encrypt and sign). We implemented random access as described in Section 6.1.

[4] The prefix `.x-md-x.` is chosen arbitrarily. SiRiUS uses it only to locate *md-files* on the remote server. No special steps are needed to deal with files with the same name as the prefix.

[5] The chunk size depends on the NFS implementation and most implementations optimize for 8192-byte blocks.

Test	File Size	Kernel NFS	DumbFS	SiRiUS
File Creation	0	0.4	3.4	14.5
File Deletion	0	0.3	0.4	1.1
Sequential Read	8 KB	0.9	1.4	18.0
Sequential Write	8 KB	1.1	2.0	21.9
Sequential Read	1 MB	96.7	97.8	223.8
Sequential Write	1 MB	100.0	102.7	632.9

Table 2. Benchmark Timings for Basic Operations. Numbers given are in milliseconds.

4.7. Caching

Our implementation avoids unnecessary operations by aggressively caching meta data and integrity information. All the caching code is implemented over the the file handle cache infrastructure. The *md-file* does not have to be fetched and verified repeatedly on a read unless we encounter an integrity or decryption failure. We can perform the same optimization for writes by the file owner. In addition, for a read operation on block i of file `foo`, we only need to compare the hash of block i to the cached file hash block of `foo`. If the hashes are the same, we do not need to fetch the hash block and verify its signature again. Unfortunately, writes update the hash block and so we are forced to perform a signature for every write. Similar optimizations are implemented for freshness verification.

4.8. Performance

We performed the tests listed in Table 2 to compare the performance of SiRiUS layered over NFS, the Linux kernel (2.4.19) NFS Client and DumbFS (a NFS pass-through proxy built using the SFS toolkit). Each test was performed on a hundred different files and the results were averaged. Our implementation uses AES-128 as the block cipher, RSA-1024 as the public key encryption algorithm, and DSA-512 as the signature scheme. The NFS server was run on a Pentium III 1.13 GHz machine and the three clients were run on a Pentium III-M 866 MHz machine.

The DumbFS benchmarks show the low overhead of using a user-level loopback server. File creation is much slower in SiRiUS because it requires key and signature generation. Deletions are slightly slower because SiRiUS has to unlink two files compared to just one for regular NFS clients. For small files, reads and writes are about 20 times slower than the kernel NFS client. For writes, the slowdown is due to the overhead of encrypting data (decrypting for reads), verifying three signatures (two for file integrity and one for freshness), and generating a signature (no signature generation for reads).

For larger files, our aggressive caching and optimizations start to pay off. SiRiUS is able to read a 1 MB file with only a $2.3\times$ slowdown, as compared with the kernel NFS. SiRiUS writes a 1 MB file with $6.3\times$ slowdown in comparison with the kernel NFS. Keeping in mind the extensive cryptographic operations involved in reads and writes, these benchmarks represent excellent performance. Reads are about $3\times$ as fast as writes because every write has to sign the modified hash block. Sequential writes on a large file can be sped up if we can cache all the changes to the file hash block before performing the signature. Unfortunately, NFSv3 does not provide a file CLOSE call which would allow this optimization.[6]

The first read (or write) of a random 8 KB block within a large file will take the same amount of time as a sequential read (or write) of an 8 KB file. Subsequent 8 KB block operations on the same file will see dramatic performance improvements (similar to those observed for large files) because of the caching and verification optimizations described in the previous section.

A profile of the current implementation shows that 63% of the time spent during a 1 MB file read is on AES decrypt. Signature generation take up about 40% of the time on a 1 MB file write. We can improve read performance by switching to a faster block cipher. The cost of signature generation for writes can be alleviated if expensive computations are performed out-of-line, as with on-line/off-line signature schemes [11]. For example, DSA signature generation is computationally expensive because it requires an exponentiation to a random number. This exponentiation is independent of the message to be signed and can therefore be precomputed by the signer before he receives the message. The SiRiUS client can reduce write times by precomputing these random exponentiations during idle cycles and using these precomputed values for signature generation when write requests are received.

5. Weaknesses

The design constraints to SiRiUS — in particular, the prohibition on modifications to the server — limit the security we can provide. We list some of these limitations here.

[6] NFSv4 supports a CLOSE call.

5.1. Deleting File System Contents

There are some inherent difficulties in securing legacy network file systems without modifying the file server. For example, since the file server cannot be trusted to perform any access controls, an attacker can break into the server and perform a denial-of-service attack by deleting all the files. SiRiUS can do very little to prevent these sorts of attacks. Without assuming anything about the underlying file system, the best that SiRiUS can do is to provide tamper-detection mechanisms. The real solution is to layer SiRiUS over a Byzantine storage system [6].

5.2. Seizing File System Control

A malicious user can replace an existing file (and its meta data) and then update the root *mdf-file* and sign it himself. In this case, unless the file reader knows who the true owner is supposed to be, file verification will succeed. The entire file system can be seized by performing this attack on all the files. However, this attack is of little significance since the file system owner will quickly detect the attack when carrying out the periodic update to the root *mdf-file*. Another solution to this problem is to publish the root *mdf-file* (which is small) on a secure server so that the freshness hash tree can be independently verified.

5.3. *d-file* Rollback

The freshness guarantees in basic SiRiUS only apply to the *md-files*. Hence, a rollback attack that replaces the newest version of the *d-file* with an older version succeeds in basic SiRiUS. Because the older version is still a legitimately-created file, this attack is not as devastating as that against *md-files* (described in Section 3.3), where revoked permissions are restored. However, *d-file* rollbacks are a problem and should be fixed.

We cannot counter this attack with the scheme for *md-file* freshness applied to *d-files* because each *d-file* can have multiple writers whereas each *md-file* has only one writer. The best solution to the *d-file* rollback problem is the one described in Section 6.4 for maintaining traditional file system semantics using union mounts.

6. Extensions

SiRiUS extensions are non-essential capabilities that improve the performance or security of basic SiRiUS. We have implemented the random access extension (for performance reasons).

6.1. Random Access and Low Bandwidth

Currently, only cryptographic file systems that operate at the block level are able to support random access and low bandwidth. As mentioned in Section 2.1, it is hard to achieve such properties in file systems that operate at a higher level. The main difficulty is in updating and verifying file integrity information without accessing the entire file.

Basic SiRiUS is easily modified to add random access and low bandwidth. We first give an overview of the random access scheme in SiRiUS. The key insight is that we can represent each file as a series of blocks, each with its own integrity information. Each file block is encrypted using a block cipher in CBC mode [10] with a different random initial value, and each block is also hashed for integrity. However, to prevent block swapping attacks, the integrity of each block needs to be related to the integrity of the entire file. A hash tree construction, similar to that used for SiRiUS meta data freshness, can be used to relate block integrity to file integrity.

Figure 5 shows the *d-file* format required for random access. The symbol $||$ denotes concatenation. The last block of the *d-file* contains all the block hashes (H_1 to H_n) and the signature (using the *FSK*) of the block hashes. Note that the last block is of variable size but is generally small.

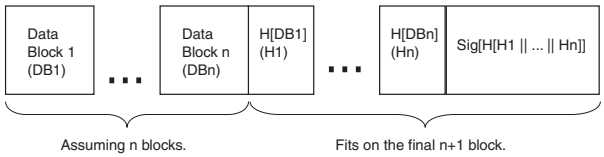

Figure 5. Data file format for random access.

We describe how random access works in SiRiUS by elaborating on the procedures for the update and the verification of file data. We assume that the block cipher used is AES in CBC mode, the hash function is SHA-1 and the signature scheme is DSA.

We first describe how a block within the *d-file* is updated. Assume that the *md-file* has been verified and we have the *FEK* and *FSK*. Also assume that we are updating block number i.

1. Encrypt the data using AES in CBC mode with the *FEK* and a random initial value.[7]
2. Hash encrypted block i and replace the hash value for block i in the final hash block. Update the new initial value for block i. Re-apply SHA-1 to the concatenation of block hashes to obtain a new file hash, and sign that with the *FSK*.

Observe that we need to fetch only two blocks to update one data block. For updates that span multiple data blocks, the procedure is applied to each affected block.

[7] The initial values for each file block can be stored in the clear with the block hash values.

Verifying a single file block j is equally easy. Fetch both the file block j and the final hash block. Hash file block j and recompute the file hash using the hashes of the other blocks in the hash block; the actual data blocks need not be retrieved. Then, verify that the signature from the hash block corresponds to the computed file hash.

6.2. Encrypted Pathnames

Adding filename encryption to SiRiUS is simple. When creating the file, use the *FEK* to encrypt the filename. Any change in the *FEK* requires the filename to be re-encrypted. When listing the contents of a directory, the SiRiUS client iterates over all the *md-files* in the directory to obtain the *FEK* of each file to which the user has access. With an *FEK*, filenames can be decrypted and displayed. Files to which the user has no access are not displayed. This list operation is potentially expensive, since two public-key operations are required for decryption and verification of the encrypted key block. We can speed up this operation by caching the file keys on the client machine.

In basic SiRiUS, directories in the file system do not have an associated *md-file*. If directory entries are encrypted, then each directory needs an associated *md-file*.

filename collisions are a potential problem because each file name is encrypted with a different key. We provide an example of the problem. Suppose that Alice wishes to rename a file `foo` to `bar`, and that, in the same directory, there are other files to which she does not have read access. There might already be another file called `bar` of which she is unaware.

The filename collision problem is solved in SiRiUS by prepending a hash of the unencrypted filename to the encrypted filename. Before a file is created or renamed, the directory can be checked for collisions. This solution comes at the expense of reducing the maximum filename length by the size of the hash.

6.3. Large-Scale Group Sharing using NNL

Basic SiRiUS is optimized for small group sharing and does not scale well when large groups of users share files. For example, when Bob revokes Alice's read access to a file, Bob must generate a new *FEK* and update all the encrypted key blocks in the *md-file* with the new *FEK*. Thus, Bob must perform n public key encryptions if there are n users sharing the file.

For large-scale sharing, we can use the NNL construction [26] for key revocation. NNL works well when the group is close to the size of the entire user set. We call this extension SiRiUS-NNL. We provide an overview of NNL's properties; we then describe the new *md-file* format and the process of file creation, sharing, and revocation.

NNL Overview. NNL [26], or Naor-Naor-Lotspiech, introduces the subset-sum framework of schemes for broadcast encryption and traitor tracing. A broadcast encryption scheme [12] is concerned with efficient transmission of a message to a group of receivers whose membership is not fixed: for example, from a satellite television provider to its subscribers. A traitor-tracing scheme [8] allows the recovery, from a box capable of decoding broadcast transmissions, of the identity of (one of) the users who colluded in the box's creation.

Broadcast encryption provides a natural model for SiRiUS' key-distribution requirements. The owner of a file will want to distribute that file's encryption key *FEK* to some set of users and its signing key *FSK* to some other set, and do so as space-efficiently as possible. Changes in permissions are equivalent to changes in the membership of one or both sets of users.

In NNL's subset-sum framework, the set of potential recipients is partitioned into subsets, each associated with a long-lived key. Each recipient belongs to a number of sets, and possesses only the keys of those sets. The broadcaster chooses a subset cover, i.e., a set of subsets whose union includes exactly those users whom he wishes to receive the transmission. He encrypts the transmission key under each included subset's key.

Basic SiRiUS is a special case of this framework. Each user belongs to one subset; that subset's long-lived key is the user's key. If there are N users in the system, of whom n should be able to decrypt (or, alternately, $r = N - n$ are revoked), then transmission size is $O(n)$, each receiver stores $O(1)$ keys.

Naor et al. provide two more sophisticated instantiations of the subset-cover framework. These are both more efficient than the basic approach when n is large (alternately, r is small). Both constructions are combinatorial, and make use of a binary tree; each user is assigned a leaf node.

In the complete-subtree instantiation, each subset is a complete subtree rooted at some node in the tree. A user is given the keys corresponding to those subtrees rooted at nodes along the path from her leaf to the root. In the complete-subtree instantiation, transmission size is $O(r \lg (N/r))$ and each receiver stores $O(\lg N)$ keys.

In the subset-difference instantiation, a subset is defined by two nodes V_i and V_j such that V_j is in the subtree whose root is V_i; the set contains all those nodes in V_i's subtree but not in V_j's subtree. A user is given the keys corresponding to those subsets that include her. In this instantiation, transmission size is $O(r)$ and each receiver stores $O(\lg^2 N)$ keys.

Halevi and Shamir [14] provide a more efficient instantiation of the subset-cover framework that has a transmission size of $O(r)$, and requires only $O(\lg^{1+\epsilon}(N))$ keys.

SiRiUS-NNL *md-file* Format. In SiRiUS-NNL, each *md-file* requires two separate NNL trees. The keys in one NNL tree are used to encrypt a file key block (in the *md-file*) which contains the *FEK* and *FSK*. Call this file key block the *FKB-write*. The keys in the other NNL tree are used to encrypt another file key block which contains only the *FEK*. Call this block the *FKB-read*.[8] With these two file key blocks, we can still separate read from read-write access, since a user with only read access will not have keys to the NNL tree for *FKB-write*. The *FKB-write* and *FKB-read* are encrypted using the appropriate choice of keys calculated using the Subset-Cover revocation algorithms (described in the NNL paper [26]) on the NNL trees.

A user's encrypted key block in the *md-file* contains the symmetric encryption keys that constitute a path from a leaf to the root of one of the two NNL trees. Since this path never changes, the encrypted key blocks are not updated on a revocation, removing the need for a public-key encryption on revocation. Since the encrypted key blocks are larger and there are additional key blocks, we are trading space for time in SiRiUS-NNL. This tradeoff is acceptable since disk space is cheap and plentiful.

File Creation, Access, and Revocation. On file creation, all the symmetric keys for two NNL trees for that file are created and stored in the owner's encrypted key block. Enough keys are created so that more users can be added in the future.

When a user is given read access to a file, the file owner obtains a set of keys from the NNL tree for *FKB-read* and creates the encrypted key block for that user with these keys. A similar procedure is used for write access.

When a user's read access to revoked, the owner regenerates the *FEK* and re-encrypts the *d-file*. She also calculates the new choice of keys for the *FKB-read* and encrypts the *FKB-read* with this new set. A similar procedure is used for write-access revocation.

6.4. Maintaining Traditional File System Semantics

In basic SiRiUS, each user owns a separate file system on the remote server. Since there is no concept of directory permissions in SiRiUS, only the owner of a file system can create or rename files on her file system. These semantics present a problem with some traditional applications such as editors. For example, Emacs creates temporary files in the working directory during editing, and, on a save, replaces the original file with the temporary copy through a rename. Many applications use this rename paradigm to guarantee atomic file updates.[9] Applications such as CVS will not work flawlessly in basic SiRiUS because non-owner users need to create new files in a repository.

Union Mounts. SiRiUS can be extended to support traditional file system semantics through the use of union mounts [30, 20]. We use a generalized union-mount system to merge SiRiUS file systems belonging to different users together to obtain a single view of the merged file systems. We refer to a SiRiUS extension with union mounts as SiRiUS-U.

The BSD union mount system allows a stack of file systems $F_1, \ldots, F_k$ to function logically as a single file system.

For each file read attempt, the stacked file systems are searched, from F_k down to F_1; the topmost one that contains a version of the file is used to answer the attempt.

Any changes to a file actually modify the copy in the topmost, read-write file system F_k. File deletions are handled by placing a "whiteout" entry in F_k that hides the lower-level file.

Union mounts are useful, for example, when compiling sources from a CD-ROM: a directory tree under `/tmp` can be stacked over the (read-only) source tree; object files and executables are created in `/tmp`, but appear side-by-side with the sources.

When a SiRiUS-U client mounts a user's file system, the client carries out a union mount of all other users' corresponding file systems on the remote file server. All the freshness files and *md-files* for all users are visible to the SiRiUS-U client so that it can locate and verify files to which it has access.

To handle name-space collisions, SiRiUS-U requires a logical ordering of the union-mounted stack of SiRiUS file systems that is different from that in previous union-mount implementations. The logical ordering required is temporal: The most recent version of a file is ordered at the top of the search stack. We illustrate this ordering with an example. Suppose Bob is the owner of file `foo`, Alice alone has write access to the file, and Carol has read access. Suppose Alice performed the last write on `foo` and Carol tries to read `foo` some time later. Carol's SiRiUS-U client checks the union-mounted file system and locates both Bob's and Alice's versions of `foo`. Since Alice's version is more recent, the SiRiUS client displays her version to Carol.

If Alice is not allowed to write to `foo`, then Carol's client should not consult Alice's version regardless of its temporal ordering. Bob's meta data files must therefore specify which users are allowed to modify a given file,

[8] These file key blocks are not present in basic SiRiUS.

[9] Emacs is intelligent enough to create temporary files in the local `/tmp` directory if the user has no write access to the working directory.

and only these users' directory trees should be consulted in accessing a file.

Using union mounts allows SiRiUS to maintain the file system semantics discussed earlier. Furthermore, SiRiUS-U retains the security of the original scheme. In fact, SiRiUS-U can also guarantee the freshness of *d-files*.

Freshness Guarantees Using Union Mounts. If union mounts are used, the *md-file* freshness scheme can be extended to *d-files* to solve the *d-file* rollback attack mentioned in Section 5.3.

Each SiRiUS user also computes a hash tree for the hash blocks of all the *d-files* that he has write access and creates a data freshness file or *df-file*. The procedures for generating the *df-files* are similar to those for generating *mdf-files*. The root *df-file* is periodically time-stamped and signed by the user.

The union mount of all user file systems ensures that all *df-files* for each user are visible to everyone. These *df-files* are used verify the freshness of file data. For example, suppose that file `foo` has two writers, Bob and Carol, and that Alice wants to read the file. The union mount shows that the latest version of `foo` was last modified by Bob, so Alice verifies the *d-file* freshness using Bob's *df-files*.

7. Related Work

The Self-Certifying File System [18] (SFS) provides authentication and channel security for accessing remote file systems. Access control in SFS relies on all file requests' passing through the trusted SFS server. The use of a trusted server in SFS removes the need for end-to-end security. In contrast, SiRiUS is designed on the assumption of an untrusted file server and therefore provides end-to-end security with cryptographic access controls.

The Cryptographic File System (CFS) [2] associates a single symmetric key with each file system. The pathnames and data of the file system contents are encrypted before being written to disk. Access control is determined by possession of the file system key. Hence, CFS allows only coarse sharing with no read-write separation or per-file sharing. In contrast, SiRiUS allows fine grained sharing on a per-file basis.

CryptFS [35] is similar to CFS except that symmetric keys are associated with groups of files. These group file keys permit group sharing but no read-write access controls. The Extended Cryptographic File System (ECFS) [1] extends CFS to provide file integrity. The Cryptographic Storage File System (CSFS) [13] is similar to CFS but also supports file integrity and group sharing of files. However, CSFS does not provide read-write access controls. CSFS also relies on a highly trusted group database server that determines group membership (and hence access control). The Transparent Cryptographic File System (TCFS) [7] is another extension of CFS. Each user in TCFS possesses a master key which is used to protect unique file keys. TCFS has a group threshold sharing scheme that require a certain number of active group members before group shared files become accessible. Apart from this scheme, no other form of sharing is supported. In contrast, SiRiUS provides per file read-write sharing and does not rely on access control servers.

Secure Network Attached Disks (SNAD) [24] is a block storage system providing end-to-end encryption and integrity of file data and meta data. Files are protected with unique symmetric keys, in turn protected with the asymmetric key of each user given access to the file. However, access permissions are stored and enforced by the remote file server. This model requires strong authentication of users along with user trust in the server to enforce access control. The Secure Untrusted Data Repository (SUNDR) [19] uses block servers to store blocks of data that clients interpret as a file system. The file system implementation resides entirely on the client. SUNDR provides end-to-end encryption using per-file keys. File integrity is implemented using hash trees. SUNDR relies on the server's enforcing access control. Both SUNDR and SNAD require new storage servers. In comparison, SiRiUS does not require new infrastructure and also allows each user to manage his own access permissions without having to trust the file server.

Plutus [15] provides end-to-end security, along with a novel cryptographic group sharing system with lazy revocation. CryptosFS [29] explores the use of public key cryptography to replace access control mechanisms. CryptoFS relies on the file server correctly to verify access by users. Furthermore, obtaining the asymmetric keys only grants access to the enciphered file; files are also encrypted using a symmetric key that users must obtain (through some out-of-band channel) to decrypt and encrypt the file contents. In contrast, SiRiUS has a simpler key management scheme with user-managed access control. The Swallow distributed file system [31] enforces access control through client side cryptography. Swallow clients encrypt their files before storing them on the remote file servers. However, Swallow does not provide fine grained file sharing. If the owner of a file allows read access to another user, that user can also modify the file undetectably.

Zero-Interaction Authentication (ZIA) [9] aims to secure mobile devices even against physical attacks. ZIA implements file access control via a cryptographic file system that communicates with a physical authentication token. Public key cryptography is used to authenticate a physical token to the mobile device. File sharing is implemented using symmetric ciphers. SiRiUS can use ZIA to secure the master keys on the client machine.

8. Summary and Conclusion

The use of SiRiUS is compelling in situations where users have no control over the file server (such as Yahoo! Briefcase or the P2P file storage provided by Farsite). We believe that SiRiUS is the most that can be done to secure an existing network file system without changing the file server or file system protocol.

Acknowledgments

The authors thank Monica Lam, Ben Pfaff, Mendel Rosenblum, and the other members of the Stanford Collective Group for their helpful discussions. The authors also thank Kevin Fu and the anonymous reviewers for their comments. Special thanks to Constantine Sapuntzakis for comments on maintaining standard file system semantics, and Robert McGrew for pointing out an attack (now fixed). The authors are grateful to Mendel Rosenblum and VMware for providing us with their software for development work.

References

[1] D. Bindel, M. Chew, and C. Wells. Extended cryptographic file system. Unpublished manuscript, December 1999.

[2] M. Blaze. A cryptographic file system for Unix. In *Proceedings of the First ACM Conference on Computer and Communications Security*, pages 9–16. ACM, November 1993.

[3] D. Boneh and M. Franklin. Identity based encryption from the Weil pairing. In J. Kilian, editor, *Proceedings of Crypto 2001*, volume 2139 of *LNCS*, pages 213–229. Springer-Verlag, August 2001.

[4] B. Callaghan. *NFS Illustrated*. Addison-Wesley, December 1999.

[5] B. Callaghan, B. Pawlowski, and P. Staubach. NFS version 3 protocol specification. RFC 1813, June 1995.

[6] M. Castro and B. Liskov. Practical Byzantine fault tolerance. In *Proceedings of the Third Symposium on Operating Systems Design and Implementation (OSDI*. USENIX, Febuary 1999.

[7] G. Cattaneo, L. Catuogno, A. D. Sorbo, and P. Persiano. The design and implementation of a transparent cryptographic file system for UNIX. In *Proceedings of USENIX Technical Conference, FREENIX Track*. USENIX, June 2001.

[8] B. Chor, A. Fiat, and M. Naor. Tracing traitors. In Y. Desmedt, editor, *Proceedings of Crypto 1994*, volume 839 of *LNCS*, pages 257–70. Springer-Verlag, August 1994.

[9] M. Corner and B. Noble. Zero-interaction authentication. In *Proceedings of the Eighth International conference on Mobile Computing and Networking (MOBICOM)*, pages 1–11. ACM, 2002.

[10] M. Dworkin. Recommendation for block cipher modes of operation. Special Publication 800-38A, NIST, 2001.

[11] S. Even, O. Goldreich, and S. Micali. On-line/off-line digital signatures. In G. Brassard, editor, *Proceedings of Crypto 1989*, volume 435 of *LNCS*, pages 263–277. Springer-Verlag, August 1989.

[12] A. Fiat and M. Naor. Broadcast encryption. In D. Stinson, editor, *Proceedings of Crypto 1993*, volume 773 of *LNCS*, pages 480–91. Springer-Verlag, August 1993.

[13] K. Fu. Group sharing and random access in cryptographic storage file systems. Master's thesis, Massachusetts Institute of Technology, June 1999.

[14] D. Halevi and A. Shamir. The LSD broadcast encryption scheme. In M. Yung, editor, *Proceedings of Crypto 2002*, volume 2442 of *LNCS*, pages 47–60. Springer-Verlag, August 2002.

[15] M. Kallahalla, E. Riedel, R. Swaminathan, Q. Wang, and K. Fu. Plutus — scalable secure file sharing on untrusted storage. In *Proceedings of the Second USENIX Conference on File and Storage Technologies (FAST)*. USENIX, March 2003.

[16] U. Maheshwari, R. Vingralek, and W. Shapiro. How to build a trusted database system on untrusted storage. In *Proceedings of the Fourth Symposium on Operating Systems Design and Implementation (OSDI)*. USENIX, October 2000.

[17] D. Mazières. A toolkit for user-level file systems. In *Proceedings of the USENIX Technical Conference*, pages 261–274. USENIX, June 2001.

[18] D. Mazières, M. Kaminsky, M. F. Kaashoek, and E. Witchel. Separating key management from file system security. In *Proceedings of the Seventeenth ACM Symposium on Operating Systems Principles (SOSP)*, pages 124–139. ACM, 1999.

[19] D. Mazières and D. Shasha. Don't trust your file server. In *Proceedings of the Eighth Workshop on Hot Topics in Operating Systems*, pages 113–118, May 2001.

[20] M. K. McKusick, K. Bostic, M. Karels, and J. Quarterman. *The Design and Implementation of the 4.4 BSD Operating System*. Addison-Wesley, 1996.

[21] R. Merkle. A digital signature based on a conventional encryption function. In C. Pomerance, editor, *Proceedings of Crypto 1987*, volume 293 of *LNCS*, pages 369–378. Springer-Verlag, August 1987.

[22] Microsoft. Common internet file system (CIFS). Located at `http://www.ubiqx.org/cifs/`.

[23] Microsoft. Federated, available, and reliable storage for an incompletely trusted environment (Farsite). Located at `http://research.microsoft.com/sn/Farsite/`.

[24] E. Miller, D. Long, W. Freeman, and B. Reed. Strong security for distributed file systems. In *Proceedings of the Twentieth IEEE International Performance, Computing, and Communications Conference*, pages 34–40, April 2001.

[25] A. Muthitacharoen, R. Morris, T. Gil, and B. Chen. Ivy: A read/write peer-to-peer file system. In *Proceedings of the Fifth USENIX Symposium on Operating Systems Design and Implementation (OSDI)*. USENIX, December 2002.

[26] D. Naor, M. Naor, and J. Lotspiech. Revocation and tracing schemes for stateless receivers. In J. Kilian, editor, *Proceedings of Crypto 2001*, volume 2139 of *LNCS*, pages 41–62. Springer-Verlag, August 2001.

[27] OceanStore Project. Located at `http://oceanstore.cs.berkeley.edu/`.

[28] OpenSSL Project. Located at `http://www.openssl.org/`.

[29] D. P. O'Shanahan. CryptosFS: Fast cryptographic secure NFS. Master's thesis, University of Dublin, 2000.

[30] J.-S. Pendry and M. K. McKusick. Union mounts in 4.4BSD-Lite. In *Proceedings of USENIX Technical Conference*, pages 25–33. USENIX, January 1995.

[31] D. Reed and L. Svobodova. Swallow: A distributed data storage system for a local network. In A. West and P. Janson, editors, *Local Networks for Computer Communications*, pages 355–373. North-Holland, 1981.

[32] R. Sandberg, D. Goldberg, S. Kleiman, D. Walsh, and B. Lyon. Design and implementation of the SUN network filesystem. In *Proceedings of the Summer USENIX Conference*, pages 119–130. USENIX, 1985.

[33] R. Srinivasan. Remote Procedure Call Protocol version 2. RFC 1813, August 1995.

[34] Yahoo! Briefcase. Located at `http://briefcase.yahoo.com/`.

[35] E. Zadok, I. Badulescu, and A. Shender. Cryptfs: A stackable vnode level encryption file system. Technical Report CUCS-021-98, Columbia University, 1998.