
Privacy analysis and enhancements for data sharing in *nix systems

Aameek Singh*

Storage Systems
IBM Almaden Research Center
650 Harry Road, San Jose, CA – 95120, USA
E-mail: aameek.singh@us.ibm.com
*Corresponding author

Ling Liu and Mustaque Ahamad

College of Computing
Georgia Institute of Technology
801 Atlantic Drive, Atlanta, GA – 30332, USA
E-mail: lingliu@cc.gatech.edu
E-mail: mustaq@cc.gatech.edu

Abstract: In this paper, we analyse the data sharing mechanisms of *nix systems and identify an immediate need for better privacy support. For example, using a simple insider attack we were able to access over 84 GB of private data at one organisation of 825 users, including 300 000 e-mails and 579 passwords to financial and other private services websites, without exploiting any technical vulnerability.

We present two solutions to address this problem:

- 1 an administrative auditing tool which can alert administrators and users when their private data is at risk
- 2 a new View Based Access Control (VBAC) mechanism which provides stronger and yet convenient privacy support.

We also describe a proof-of-concept filesystem-based implementation and performance analysis of VBAC. Our evaluations with three well-known filesystem benchmarks show little overhead of using VBAC.

Keywords: unix privacy; private data sharing; access control; View Based Access Control; VBAC.

Reference to this paper should be made as follows: Singh, A., Liu, L. and Ahamad, M. (2008) 'Privacy analysis and enhancements for data sharing in *nix systems', *Int. J. Information and Computer Security*, Vol. 2, No. 4, pp.376–410.

Biographical notes: Aameek Singh is a Research Staff Member in the Storage Systems group at IBM Almaden Research Center. Singh graduated with a PhD from the College of Computing at Georgia Institute of Technology in 2007. His research interests are in the area of storage systems and data intensive distributed systems. His PhD dissertation proposed new access control and storage management techniques for enterprise storage-as-a-service environments. Singh has over 20 refereed publications in international journals and conferences and is a co-inventor on more than ten patent applications. He was a recipient of the IBM PhD fellowship for 2005–2006 and 2006–2007.

Ling Liu is an Associate Professor in the College of Computing at Georgia Tech. There, she directs the research programme in Distributed Data Intensive Systems Lab examining research issues and technical challenges in building distributed computing systems that can grow without limits. She has published over 200 international journal and conference articles. She is an internationally recognised expert in the areas of Database Systems, Distributed Systems, Internet Systems, and Web Services. She is currently on the editorial board of several top international journals, including *IEEE TKDE*, *IEEE Transaction on Service Computing (TSC)* and *VLDB Journal*.

Mustaque Ahamad is a Professor in the School of Computer Science at the College of Computing at Georgia Tech and Director of the Georgia Tech Information Security Center, a National Security Agency (NSA) Center of Excellence in Information Assurance Education. As Director of an interdisciplinary centre uniquely focused on usable security, Ahamad is responsible for the integration of research that empowers everyday users to better protect themselves and take charge of their online security and privacy. Ahamad is a leader in developing effective solutions to safeguard personal digital information against current cyber security threats, such as phishing, spoofing, and identity theft.

1 Introduction

Unix, Linux and its various other open-source flavours (together called *nix) are steadily growing in mainstream popularity and many enterprises now opt to set up their intranets using such *nix systems. One of the important features of these *nix systems is their ingrained multi-user support. In contrast to common single-user PC systems, these systems allow simultaneous multiple users and provide seamless mechanisms to share data between different users. For example, a user *alice* can set appropriate access ‘permissions’ on the data she wants to share with her group *students* by executing a simple `chmod` command (Linux Manual pages, 2008).

In this paper, we critically analyse the data sharing mechanisms of *nix systems. Access to shared data in these systems is dictated by the *nix *access control model*, which typically follows the original UNIX access model (Ritchie and Thompson, 1974). We aim to analyse the *privacy* support in its data sharing mechanisms. For example, how does the system assist a user to share data only with desired users and prevent private information from being leaked to unauthorised users? As part of the analysis, we also need to look at the *convenience* of using these data sharing mechanisms. This is important since lack of convenience typically compels users to compromise, intentionally or not, their security requirements to conveniently fit the specifications of the underlying access control model. Please note that we use the phrase ‘*private sharing*’ to indicate the desire of sharing data only with a select set of authorised users.

As part of our study, we analysed how users use access control for their data sharing needs in practice. We conducted experiments at two *nix organisations of many hundred computer-literate users each. Surprisingly, we found that large amount of private data was accessible to unauthorised users. In many cases, the user’s definition of an ‘authorised user’ does not match the underlying system’s definition, that leads to such a breach. This observation is best exemplified in the following scenario. Many users

attempt to privately share data by using execute-only permissions for their home directories (complete *nix access permissions are discussed later in Section 2). This prevents other users from listing the contents of the directory, but they can `cd` into it and any user who *knows* the name of a subfile/directory can access it with appropriate permissions. The data owner *authorises* some users by explicitly giving them the names of the subfile/directory through out-of-band mechanisms like personal communication or e-mail. However, this authorisation is not the same as the system's authorisation. From the underlying system perspective, it is assumed that any user who issues the command with the right file/directory name is authorised. Thus, users who simply guess the subfile/directory name can also access the data. Along with that, setting execute-only permissions on the home directory to share *one* subdirectory, puts all other subdirectories (sibling to the shared directory) also at risk, which if not protected appropriately, can be accessed.

We were able to effectively exploit these shortcomings in our studies. At one organisation of 825 users, over 84 GB of data was accessible, including more than 300 000 e-mails and 579 passwords to financial websites like *bankofamerica.com* and other private websites like medical insurance records. Importantly, the attack does not always need to *guess* directory names, but can find actual names from unprotected command history files (*.history*, *.bash history*) or standard application directory names (*mozilla*). The reason for this surprisingly large privacy breach without exploiting technical vulnerabilities like buffer overflows or gaining elevated privileges, is the combination of lack of system support and user or even applications' privacy-indifferent behaviour either mistakenly or for lack of anything better. Also, as we discuss later in Section 2, even for an extremely privacy-conscious user with all available tools, it is tough to protect private data in many situations.

The attack described in this paper is a form of an *insider* attack, in which the attacker is inside the organisation. The attacker could be a disgruntled employee, contractor or simply a curious employee trying to access the salaries chart in the boss's home directory. According to a recent study by the US Secret Service and CERT (Carnegie Mellon Software Engineering Institute, 2008), such attacks are on a rise with 29% of the surveyed companies reporting having experienced an insider attack in the past year (Cappelli and Keeney, 2008) (it is usually believed that such attacks are much under-reported for lack of concrete evidence or fear of negative publicity (United States Secret Service and CERT Coordination Center, 2008)). Also, in complete congruence to our attack, the report finds that:

“Most incidents were not technically sophisticated or complex – that is they typically involve exploitation of non-technical vulnerabilities.”

We propose two solutions to address this problem:

- 1 A Privacy Auditing Tool that analyses the '*privacy health*' of an enterprise. Such auditing can be combined with the periodic virus scans and other security audits regularly employed by enterprises.
- 2 A new access control model which provides stronger privacy protection and yet convenient data sharing mechanisms. The View-Based Access Control (VBAC) model allows data owners' to define views of their data that can be seen by other users, while protecting other data from unauthorised users. We provide a file system implementation of VBAC and show that the overheads incurred by it are minimal.

An interesting feature of our VBAC implementation is our highly *data* oriented focus. There has been a significant amount of work on *nix security which has been primarily focused on preventing *processes* from gaining elevated privileges or sandboxing *processes* from accessing certain data (Wagner, 1999; Jaeger and Prakash, 1994; Goldberg *et al.*, 1992; Linux Manual pages, 2008). In contrast, our approach focuses on private *data* that needs to be protected. Our implementation allows administrators to use our approach only for data that is considered private – for example, user home directories and thus, does not incur any overheads for other data in the system. To summarise, this paper makes four unique *contributions*:

- 1 *Privacy analysis of *nix systems*: To the best of our knowledge, this is the first work that evaluates the privacy characteristics of real multiuser *nix installations. We analyse the *nix access control from a privacy perspective and present results on how users share their data and how much private information can be accessed by unauthorised users in real *nix installations.
- 2 *Data sharing principles*: We identify a number of useful design principles that allow for private and convenient data sharing in multi-user environments. We also analyse existing access control tools like *nix permissions model, POSIX Access Control Lists (ACLs) (Grunbacher and Nuremberg, 2008), umask (Linux Manual pages, 2008) on these principles.
- 3 *Privacy enhancements*: We present two approaches that can enhance the privacy characteristics of *nix systems. The first approach is a conservative approach that only measures the privacy health of the system and alerts users of potential threats. The second approach is more proactive and utilises VBAC to provide safer data sharing. We also describe an implementation of VBAC and its performance analysis.
- 4 *User education*: The privacy analysis and data sharing principles presented in this paper help us in devising various privacy enhancement approaches. More importantly they bring into light a massive threat to user privacy in current systems, which can be exploited overnight by a few hundred lines of code. This work thus contributes directly to increasing user awareness and education in protecting private data.

The rest of the paper is organised as follows. In Section 2, we discuss various issues that lead to privacy breaches. We also present case studies at two *nix installations which demonstrate these breaches. In Section 3, we describe a number of useful privacy principles and analyse existing tools. Section 4 gives an overview of the two privacy enhancement approaches. Section 5 presents a detailed discussion of the VBAC access control model including its design goals, implementation and performance evaluation. We describe related work in Section 6 and conclude in Section 7.

2 Data privacy: vulnerability analysis

In this section, we discuss various scenarios that lead to privacy breaches and present the results of our case studies. This analysis will provide us insights into desirable privacy and convenience characteristics. We start off with a description of *nix access control, to establish a common reference model for subsequent discussions.

2.1 *nix access control

The *nix access control primarily follows from the original UNIX access control model (Ritchie and Thompson, 1974; Ritchie, 1978). In this discretionary access model, each file system object (file, directory, links) has an associated owner who controls the access to that object. This access can be granted to three kinds of users:

- 1 owner
- 2 group
- 3 others.

The *owner* is the object owner, the *group* is a set of users to which the owner belongs (for example, a user group for students, faculty) and *others* are all other users (all users except the owner and the group). Also, note that the *owner* permissions take precedence over *group*, that is, the owner will have access as dictated by the *owner* permission bits, even though he/she also belongs to the *group*. Further, the granted access is of three types:

1 Read

For a file, this means that a user can read a file. For a directory, this means that a user can list its contents using `ls` (Linux Manual pages, 2008). For links, the permissions are for the object that are pointed-to by the link and the link's permissions itself are not used. The read permission is represented by a 'r' or a numeric value of 4.

2 Write

For a file, the write permission allows a user to write to it. For a directory, it allows a user to create, remove or rename directory contents. For links, permissions are again for the pointed-to object. The write permission is represented by a 'w' or a numeric value of 2.

3 eXecute

For a file, the execute permissions allows running the file as a program (for example, a shell script or a compiled C program). For directories, it allows users to traverse that directory and if the directory contents have appropriate permissions, the user can then access those contents. For example, to access directory `grand-child` with path `dir/child/grand-child`, both `dir` and `child` need to have execute permissions. The execute permission is represented by a 'x' or a numeric value of 1.

The permissions bits are listed in the '*owner, group, others*' order. Table 1 shows an example directory with its 3×3 access control matrix. It will be listed as `rwX r-X -X` or in numeric terms `751` ($7 = r(4) + w(2) + x(1)$ and so on). For that directory, the owner can list its contents, create, remove, or rename its children and can traverse into that directory. The users belonging to the group of the owner can list its contents and traverse into the directory but not modify its contents. The *other* users can only traverse into the directory and not list or modify its contents.

Table 1 Example directory permissions

<i>Permission</i>	<i>Owner</i>	<i>Group</i>	<i>Others</i>
Read	X	X	
Write	X		
eXecute	X	X	X

To improve the granularity of user permissions, lately POSIX ACLs (Grunbacher and Nuremberg, 2008) have been introduced into many *nix variants. They allow setting permissions at individual user levels as opposed to *group* or *others*. However, the {*owner, group, others*} is still the most dominant paradigm and many *nix installations do not even have ACLs enabled. We will analyse the ACLs later in Section 3.

Additionally, a user-specific `umask` (Linux Manual pages, 2008) setting is used to decide permissions for a newly created file. The command specifies a mask for permissions to be *disallowed* for the owner, group or other users. For example, to always create a new file with `rwx` permissions for *owner* and no permissions for *group* or *others* (that is, permission 700), the command ‘`umask 077`’ can be stored in the user initialisation file (like `.profile`).

It is possible to override the *nix access control at a file system level. For example, the distributed Andrew File System (AFS) (Satyanarayanan, 1990) has a different set of permissions settings and thus data residing on that file system can use its settings. In fact our privacy enhancing VBAC approach also modifies the access control at that layer. We discuss this later in Section 5.

2.2 Privacy breaches

In this section, we discuss various privacy breaches that occur in *nix systems. We also present results of our study at two *nix installations that demonstrate the insufficiency of current access control mechanisms. The discussion below primarily explores the popular {*owner, group, others*} *nix paradigm. The advanced mechanisms like ACLs enhance privacy in only a few cases and we will demonstrate a clear need of a new, more complete solution.

2.2.1 Selective data sharing

The first kind of privacy breach occurs due to the need to ‘*selectively*’ share data. The selectivity can be of two kinds:

1 Data selectivity

Data selectivity is when a user wants to share only a few (say one) of the subdirectories in the home directory. So, an authorised user is allowed to access only the shared subdirectory, but not any of the sibling directories. In order to do this correctly, the owner needs to follow two steps:

- Step 1 set appropriate permissions to the shared subdirectory (at least execute permissions on the entire path to the subdirectory and the sharing permissions on the subdirectory)

Step 2 *remove permissions from the sibling subdirectories.*

The second step is unintuitive, since the user needs to act on secondary objects that are not the focus of the transaction. Also, any new file being created needs to be protected.

2 User selectivity

In many situations, users need to share data with an *ad hoc* set of users that do not belong to a single user group or are only a subset of a user group, and are not the entire *others* set. In this situation, the permissions for *group* or *others* are not sufficient. Also, creating a new user group requires administrative assistance which is not feasible in all cases.

In order to do selective data sharing, currently owners mostly use execute-only permissions on the home directories. The perception is that since users cannot list the contents of the directory, they cannot go any further than traversing into the home directory unless they know the exact name of the subdirectory. Now, the owner can authorise desired users by giving them the names of the appropriate subdirectories. Those *authorised* users can traverse into the home directory and then use the subdirectory name to `cd` into it (without having to list the contents of the home directory). From data selectivity perspective, it is assumed that they cannot access the rest of the contents and from user selectivity perspective, unauthorised users cannot access any contents.

However, the underlying system cannot distinguish between such authorised or unauthorised users. Any user who can *guess* the subdirectory name can actually access the data. For an attacker inside the organisation, this is not a Herculean task. For example, for a computer science graduate school, it is highly likely that users will have directories named *research*, *classes* or *thesis*. An easy way of creating such a list of names is by collecting names from users that actually have *read* permissions on the home directories. Within the context of a single organisation, or in general human psychology, it is likely that many users have similar directory names. This is essentially a form of *social engineering* (Mitnick *et al.*, 2002) in which users and not systems are manipulated to reveal confidential information (a well-know hacker, Kevin Mitnick said “... social engineering was extremely effective in reaching my goals without resorting to using a technical exploit ...” (Slashdot, 2008)). Of course, one simple solution is to use cryptic directory names unlikely to be guessed (security-by-obfuscation). The problem then becomes similar to the fundamental problem with passwords. Keeping commonly used passwords means that they can be guessed and cryptic passwords are tough to remember. However, the problem is more severe, since while a single password is enough to access thousands of files, a user cannot be expected to keep cryptic names for tens or hundreds of files and directories.

Secondly, many times directory names do not need to be guessed at all. The names can be extracted from *history* files (like *.history* or *.bash history*), that contain the commands last executed by the owner, like `cd`, which will include real directory names. In fact, in our experiments we found around 20%–30% of all users had readable history files and around 40% of the total leaked data was obtained from the directory names extracted from these history files.

Thirdly, it is not always user created directories that leak information. Many *applications* use standard directory names and fail to protect critical information. For example, the famous Mozilla web browser (Mozilla, 2008) stores the profile directory in

~/ .mozilla and had that directory world-readable (Mozilla Bug Report, 2008) in many cases, till as late as 2003 (the Mozilla project was initiated in 1994). Many *nix installations with the browser installed before that have this vulnerability and we were able to obtain around 575 password to financial and private websites (because users saved passwords without encrypting them). In addition, their browser caches, bookmarks, cookies and histories were also available. The browser Opera (2008) also has a similar vulnerability, though to a lesser extent. While it can be argued that it is the responsibility of application developers to ensure that this does not happen, we believe that the underlying system can assist users and applications in a more proactive manner.

The POSIX ACLs (Grubacher and Nuremberg, 2008), if used help in achieving only user selectivity. They do not address the data selectivity requirements or prevent leaking of application data.

2.2.2 Metadata privacy

So far, we have only talked about the privacy breach for file *data*. However, there are many situations in which users are interested in protecting even the metadata of the files. The metadata contains information like ownership, access time, updation time, creation time and file sizes. There are scenarios where a user might obtain confidential information by just looking at the metadata. For example, an employee might be interested in knowing how big is his annual review letter or did the boss update it after the argument he had with her?

The *nix access control does not provide good metadata privacy. Even if users only have execute permissions on a directory, as long as they can guess the name of the contained file, its metadata can be accessed even if the file itself does not have any read, write or execute permissions on it. Thus, if a user has to share even a single file/directory within the home directory (thus, requiring at least execute permissions), all other files contained in the home directory have lost their metadata privacy.

Of course, a solution is to put such files in a separate directory and protect them. However, in many cases these files might be accessed by standardised applications making it infeasible to move (for example, how to protect metadata privacy of shell initialisation files (*.profile*), or history files, which are always created in the home directory). Also, from our experience, many users like to keep their active files in the home directory itself and find it inconvenient to have a deep directory structure.

Again, lack of convenient support from the system leads to privacy breaches, in this case leaking file metadata.

2.2.3 Data sharing convenience

User convenience is an important feature of an access control implementation. If users find it tough to implement their security requirements, they are likely to compromise the security requirements to easily fit the underlying access control model. This can be seen as one of the reasons why encryption file systems are not in widespread use, even though they guarantee maximum security.

From our analysis of the *nix access control, along with some of the issues discussed earlier, we found the following two data sharing scenarios in which there is no convenient support for privacy:

1 Sharing a deep-rooted directory

For a user to share a directory that is multiple levels in depth from the home directory, there needs to be at least execute permissions on all directories in the path. This in itself:

- leaks the path information
- puts sibling directories at risk
- leaks metadata information for sibling directories.

In order to prevent this, since most operating systems do not allow hard links to directories anymore, a user would have to create a new copy of the data. And since users are more careless with permissions for deep rooted directories (they protect a higher level directory and that automatically protects children directories), a copy of such a directory could have privacy-compromising permissions.

2 Representation of shared data

In many circumstances the way one user represents data might not be the most suitable way for another user. For example, while an employee might keep his resume in a directory named `job-search`, it is clearly not the most apt name to share with his boss. The employee might want her to see the directory simply as `CV`. Changing the name to meet the needs of other users is not an ideal solution. This again shows the lack of adequate system support for private and convenient data sharing.

It is important to recognise that even an extremely privacy-conscious user cannot protect data at all times. Exhaustive user efforts to maintain appropriate permissions on all user and *application created* data will still be insufficient to protect metadata privacy or allow private sharing of deep rooted directories with user-specific representation.

2.6 Case studies

As part of our study, we conducted experiments at two geographically and organisationally distinct *nix installations. Users at both installations (CS graduate schools) are highly computer literate and can be expected to be familiar with all available access control tools.

For our analysis, we consider the following data to be private:

- All user e-mails are considered private.
- All data under an execute-only home directory is considered private.
- Browser profile data (including saved passwords, caches, browsing history, cookies) is considered private.

The second assumption above merits further justification. It can be argued that not every subdirectory under an execute-only home directory is meant to be private (for example, a directory named *public*). However, we believe our definition to be a practical one. The ideal measurement would require active user participation in the study who, by anecdotal experience, once told of the threat immediately removed all permissions from their home directories. Also, the semantics of the execute-only permission set dictate

that any user other than the owner cannot list the contents of the directory and since the owner never *broadcasts* the names of the shared directories, an unauthorised user *should not* be able to access that data. And since we do not include in our measurements any obviously-private data from home directories of users with *read* permissions (for example, world-readable directories named `personal`, or `private`), we believe the two effects to approximately cancel out.

2.7 Modus operandi

Next, we describe the design of our attack that scans user directories and measures the amount of private data accessible to unauthorised users. This discussion is also important since the design is eventually used to develop an auditing tool discussed later. The attack works in multiple phases. The first step is to obtain directory name lists which can be tried against users with execute-only home directories. Three strategies are used to obtain these lists:

- 1 *Static lists*: These are manually entered names of directories likely to be found in the context of the organisations – CS graduate schools. For example, ‘*research*’, ‘*classes*’, ‘*papers*’, ‘*private*’ and their variants in case (‘*Research*’) or abbreviations (‘*pvt*’).
- 2 *Global lists*: These lists are generated by obtaining the directory names from home directories of users that have *read* permissions.
- 3 *History lists*: These are user specific lists generated by parsing users’ history files, if readable. We used a simple mechanism, parsing only `cd` commands with directory names. It is possible to do more by parsing text editor commands (like `vim`) or copy/move commands.

In the next step the tool starts a multi-threaded scanning operation that attempts to scan each user directory. For users with *no* permissions, no scanning is possible. For users with *read* permissions, as discussed earlier, since there is no precise way of guessing which data would be private, we only measure e-mail and browser profile statistics. Finally, for users with *execute-only* permissions, along with e-mail and browser profile statistics, we also attempt to extract as much data as possible using the directory name lists prepared in the first step.

2.7.1 Evaluating e-mail statistics

This is done by attempting to read data from standard mailbox names – ‘*mail*’, ‘*Mail*’, ‘*mbox*’ in the home directory and the mail inboxes in `/var/mail/userName`. A `grep` (Linux Manual pages, 2008) like tool is used to measure:

- number of readable e-mails
- number of times the word ‘*password*’ or its variants appeared in the e-mails.

2.7.2 Evaluating execute-only data statistics

For users with execute-only permissions on the home directory, the scanner uses the combination of static, global and the user's history lists to access possible subdirectories. Double counting is avoided by ensuring that a name appearing in more than one list is accounted for only once and by not traversing any symbolic links. While scanning the files, counts are obtained for the total number of files and the total size of the data that could be accessed.

2.7.3 Evaluating browser statistics

The mozilla browser (Mozilla, 2008) stores user profiles in the `~/ .mozilla` directory. This directory used to be world-readable till as late as 2003 when the bug was corrected (Mozilla Bug Report, 2008). Within that profile directory, there are subdirectories for each profile that has been used by that user. The default profile is usually named `'default'` or `'Default User'`. So even in case the `.mozilla` directory had execute-only permissions, it is possible to access default profile directories (unless a user specifically removed permissions). Within the profile directory, there is another directory with a randomised name ending in `' .sit'`. Since the parent directories had `read` permissions, the randomisation provides no security and the name is visible. Within this directory, the following files exist and (with this bug) were readable:

- Password database

A file with the name of type `'12345678.s'`. This contains user logins and passwords saved by mozilla when the user chooses to save them. Ideally, users should use a cryptographic master key to encrypt these passwords, but as our results will show many users do not encrypt their passwords. For such cases, mozilla stores the passwords in a base-64 encoding (indicated by the line starting with a `~` in the passwords file), which can be trivially decoded to get plaintext passwords.
- Cookies

The `cookies.txt` file contains all browser cookies. Many websites including popular e-mail services like Gmail (Google Mail, 2008), Hotmail (Microsoft Hotmail, 2008) allow users to automatically login by keeping their usernames and passwords (encrypted) in the cookies file. Hijacking this cookie can allow a malicious user to login into these accounts. For many other cookies related attacks (see Sit and Fu, 2001).
- Cache

This is a subdirectory that contains the cached web pages visited by the user.
- History database

Web surfing history, which many sophisticated viruses and spyware invest resources to collect, are also readable.
- Forms database

Mozilla allows users to save their form data, stored in a file of type `'23456789.w'`, that can be automatically filled. This could include credit card numbers, social security numbers and other potentially sensitive information. Here again, users should use a master key to encrypt this information.

It is important to emphasise the reasoning behind illustrating this privacy breach. We recognise that this specific vulnerability has been patched now. Our intention is to point out that relying on application developers to provide privacy is not a perfect solution for two reasons – (1) bugs like the one exploited in this case occur many times, and (2) often application developers cannot account for all privacy implications of their design decisions, for example, the SUN Java compiler does not ensure that the compiled byte code has the same permissions as the source Java file and as byte code can be decompiled back into original source (Jad: JAVa Decompiler, 2008) this clearly has a privacy implication. Instead we contend that as an alternative, the underlying system can contribute to provide better privacy support. For example, even with this vulnerability, our viewfs based approach described later would have prevented such a breach.

2.8 Results

The complete characteristics of the two organisations are shown in Table 2. Both organisations are computer science graduate schools at two different geographical locations within the United States. At both the organisations, a significant number of users (69% and 77%) used execute-only permissions on their home directories.

Table 2 Case study organisation characteristics

<i>Org.</i>	<i># Users</i>	<i># ReadX</i>	<i># NoPerms</i>	<i># X-only</i>
Org-1	825	198	54	573
Org-2	768	136	39	593

Notes: # ReadX is the number of users with read and execute permissions to their home directories.
 # NoPerms are users with no permissions.
 # X-only are the users with only execute permissions.

Table 3 lists the amount of data extracted from execute-only home directories at Org-1 and 2.

Table 3 Data extracted from X-only home directory permissions

<i>Org.</i>	<i># Hit users</i>	<i># Hits</i>	<i># Files</i>	<i>Data size</i>
Org-1	462	2409	983 086	82 GB
Org-2	380	911	364 932	25 GB

Notes: # Hit users is the number of users that leaked private information.
 # Hits is the total number of directory name hits against all X-only users.
 # Files is the number of leaked files.
 Data-size is the total size of those files.

As can be seen, a large fraction of users indeed leaked private information – 56% and 49% of total users respectively. Recall that we do not extract any data from users with *read* permissions on their home directories; so a more useful number is the fraction of X-only users that revealed private information. That number is 80% and 64% respectively. Also, on an average, 2127 files and 177 MB of data is leaked in the first organisation for each X-only user and 960 files and 65 MB of data is leaked in the second

organisation. A partial reason for the lower numbers in the second organisation could be the fewer number of users with *read* permissions, which would have impacted the global name lists creation. Overall, we believe this to be a very significant privacy breach.

As mentioned earlier, many times the names of the subdirectories do not need to be guessed and can be obtained from the history files in the user home directories. Table 4 lists the success rate of the attack in exploiting history files. As it shows, around 40% of X-only users had readable history files which led to 40%–50% of total leaked data in size.

Table 4 Exploiting history files

<i>Org.</i>	<i># History hits</i>	<i># Files</i>	<i>Data size</i>
Org-1	253	561 254	35 GB
Org-2	237	155 826	14 GB

Notes: # History hits is the number of users with readable history files.

Files is the number of private files leaked due to directory names obtained from history files.

Data-size is the size of the leaked data.

2.8.1 E-mail statistics

Table 5 presents the results of the e-mail data extracted from users in both organisations. Recall that this data is obtained for both X-only users and the users with *read* permissions on their home directories.

Table 5 E-mail statistics

<i>Org.</i>	<i># Folders</i>	<i># E-mails</i>	<i>Size</i>	<i># Password</i>
Org-1	2509	315 919	4.2 GB	6352
Org-2	505	38 206	120 MB	237

Notes: # Folders is the number of leaked e-mail folders.

E-mails is the total number of leaked e-mails.

Size is the size of leaked data.

Password is the number of times the word '*password*' or its variants appeared in the e-mails.

As can be seen, a large number of e-mails are accessible to unauthorised users (especially at Org-1). Also, the number of times the word 'password' or its variants appear in these e-mails is alarming. Even though we understand that some of these occurrences might not be accompanied by actual passwords, by personal experience, distributing passwords via e-mails is by no means an uncommon event.

2.8.2 Browser statistics

The second organisation did not have the mozilla vulnerability since they had a more recent version of the browser installed, by which time the bug had been corrected. So the results shown in Table 6 have been obtained only from the first organisation. Looking at the results, the amount of accessible private information is enormous.

Figure 1 contains a sample of the websites that had their passwords extractable and clearly most of these websites are extremely sensitive and a privacy breach of this sort is completely unacceptable. As an interesting side statistic, out of the 579 passwords, only 308 were unique passwords, that is, 36% of passwords were repeated. This indicates that users repeat their passwords, a common habit by anecdotal experience. Also as seen from Figure 1, some obtained passwords were for accounts in other institutions and a few of them are likely to be *nix systems. Thus, it is conceivable that this password extraction can be used to *expand to other *nix installations* and thus be much more severe in scope than a single installation.

Table 6 Browser statistics at Org-1

# Users with accessible .mozilla	311
# Users with readable password DB	149
# Passwords retrievable	579
# Users with readable cookies DB	207
# Cookies retrievable	19 456
# Users with accessible caches	233
# Cached entries	20 907
# Users with readable browsing histories	256
# URLs in history	130 503

Figure 1 Sample accounts with retrievable passwords

<p>Financial websites</p> <p>www.paypal.com www.ameritrade.com www.bankofamerica.com</p>	<p>Personal websites</p> <p>adultfriendfinder.com www.hthstudents.com www.icers911.org</p>
<p>E-mail accounts</p> <p>mail.lycos.com my.screenname.aol.com webmail.bellsouth.net</p>	<p>Other institutions</p> <p>cvpr.cs.toronto.edu e8.cvl.iis.u-tokyo.ac.jp systems.cs.colorado.edu</p>

2.8.3 Miscellaneous statistics

Among few other applications at Org-1, 17 users had their Opera (2008) browser’s cookies file readable and 497 users had their e-mail address books, used by the Pine e-mail client (Pine, 2008) and stored in `~/addressbook` readable. 18 308 e-mail addresses could be obtained from these address books which can be potentially used for highly targeted spam.

2.8.4 Precautions for the privacy study

A study that evaluates data and user privacy characteristics in real systems can potentially raise certain ethical issues on whether such experiments should be conducted and potential threats reported to a public forum. We appreciate any user concern in this regard.

First, we carefully crafted our experiments to prevent any privacy violation from the study itself. Our scanning tool does not store any user-specific information on disk. When the tool runs, it obtains the list of users from `/etc/passwd`, anonymises the users and randomises the order of scanning. This order is never written out to disk or printed on screen and is purged off the memory once the tool terminates. In case of measuring the number of emails that are readable, a `grep` like tool is used to only measure the number of times the word ‘Subject’: appears at the start of a line. Also, for measuring the number of passwords that can be extracted (since Mozilla stores them in base-64 encoding) only the number of such encoded passwords is calculated without decoding the base-64 encoding. This way, we ensure that we do not violate any user privacy, though such an attack is very feasible and a malicious user can obtain user specific private data. The US Secret Service and CERT study proves that such malicious users exist on the inside and trusting internal users to be ethical is incorrect.

Secondly, we believe that this study was essential to create user awareness on this very important issue and expose existing threats to user privacy that can be exploited without significant technical sophistication. Also, our proposed defences can solve or at least mitigate these privacy protection problems immediately.

2.9 Attack severity

It is important to highlight the severity of this attack:

- Low technical sophistication

The attack is extremely low-tech; the commands used in a manual attack would be `cd`, `ls` and such. This aspect makes the threat significantly more dangerous than most other vulnerabilities.

- Low detection possibility

A version of the attack that targets only a few users a day and thus keeps overall disk activity normal has a very low probability of detection. Typical *nix installations do not keep extensive user activity logs and it is highly likely that such an attack will go unnoticed. Even if an individual user notices an unusual last-access time on one of the files, without extensive logging, it is impossible to pin point the perpetrator.

- No quick fix

Unlike other security vulnerabilities like buffer overflows, this attack uses a *design* shortcoming combined with user/application carelessness and no patches would correct this problem overnight.

- High success rate

It is important to notice that the attack had a high success rate at installations where most users are computer literate. With increasing mainstream penetration of *nix systems, most users in the future would be ordinary users who cannot be expected to fully understand the vulnerabilities. This makes this attack a very potent threat.

3 Private-sharing principles

The results presented in the previous section clearly establish the fact that there needs to be much better privacy protection in *nix installations. The possible solutions might lie in user education, use of new access control tools like ACLs (Grunbacher and Nuremberg, 2008), a new access control mechanism or possibly a combination of all three. In this section, we will concretely describe important features required for *good* privacy protection and try to identify possible solutions.

Principle 1 Do not risk more than you need to

This principle implies that sharing one directory should not endanger the privacy of data in other unrelated directories. A majority of privacy breaches occurred due to users needing to have execute permissions on their home directory to share one subdirectory, but failing to adequately protect the sibling subdirectories. We believe that while sharing some data, a user should not be expected to remember if there is any private data already existing in some sibling subdirectory. Similarly, while creating new data, the user should not be expected to remember what is being shared and how that could effect the new data. It is unintuitive and likely to fail.

Principle 2 Do not trust applications completely

The second principle dictates that protecting privacy of application created data should not be left entirely on the application developer. Even for a popular project like Mozilla, it took many years to notice and correct its data leak. Secondly, many applications might simply fail to foresee a privacy impact. For example, knowing that Java byte code can be decompiled back into source (Jad: JAva Decompiler, 2008), the Java compiler should ensure that its generated *.class* file has at least the same permissions as the source (for example, the Sun javac compiler does not ensure this). Such seemingly unrelated design goals can be easily missed by application developers.

Principle 3 Increase granularity of protection

The third principle advises to increase the granularity of data as well as users in private sharing. From users perspective, it should be possible to share data with a few users without requiring them to be a part of a certain user group (or be the complete *others* set). The new ACLs (Grunbacher and Nuremberg, 2008) do provide this facility; however, they are under-used in actual installations primarily due to lack of user education. Secondly, from the data perspective it should be possible to protect even the metadata of files, if desired.

Principle 4 Convenience, convenience, convenience

The fourth principle emphasises the need for user convenience in private data sharing. Lack of convenience, for example, while sharing a deep rooted directory or inability to represent the shared data differently for the user, will inevitably lead to improper permissions on certain data and risking privacy compromise. The *nix access control model, while admittedly not inconvenient in ordinary situations, fails to facilitate slightly involved situations like the ones described above.

Principle 5 Monitor and remind users

Finally, due to the unavoidable and error-prone human element, there should be mechanisms that can monitor the system and get a measure of the *privacy health*. We believe that this should be as ingrained into auditing tools as virus scans and boot-up password enforcement, as done by many enterprise security applications (Symantec Security Manager, 2008).

3.1 Evaluating existing mechanisms

Before we present our proposed solutions in the next section, let us analyse existing access control tools on the principles described above. For the purpose of this discussion, we consider the following *nix access control mechanisms. Note that we restrict our analysis to only discretionary access control models, as are most popular in multi-user *nix installations:

- *Baseline *nix (BASE)*: The baseline *nix access control (*owner, group, others*) model as described in Section 2.1.
- *ACLs (ACL)*: ACLs (Grunbacher and Nuremberg, 2008) allow users to specify permissions at the granularity of an individual user. However, the types of permissions are the same *read, write, execute* as the baseline *nix model.
- *umask 077 (UMASK)*: The *umask 077* (Linux Manual pages, 2008) will ensure that any new file created by a user or an application will have no permissions for users other than the owner. This can be seen as a potential way of preventing inadvertent leaking of private information.
- *Richer Access Control Semantics like AFS (RACS)*: The AFS (Satyanarayanan, 1990) has additional privileges like looking up a directory, deleting files, creating new files. Using it as an example, we intend to evaluate if making the access control richer could be the solution.

Table 7 contains our analysis of these four access control mechanisms for the private sharing principles described above with X indicating lack of support, ↑ indicating partial support and ↑↑ indicating good support.

Table 7 Analysing tools for principles

<i>Principle #</i>	<i>BASE</i>	<i>ACL</i>	<i>UMASK</i>	<i>RACS</i>
1	X	X	↑	X
2	X	X	↑	X
3	X	↑	X	↑↑
4	X	X	X	X
5	X	X	X	X

As described earlier, the baseline *nix model fails on all these principles. The ACLs provide a higher level of granularity as far as users are concerned, but they do not provide data level granularity like metadata privacy. They also fail on other principles like protecting application data, or convenient means of sharing deep-rooted directories. The *umask 077* provides partial support in preventing against inadvertent exposure of user

or application files by ensuring that when the files are created, they have no readable permissions for non-owners. However, in case the user ever modifies the permissions to share anything (even temporarily), it provides no further assistance. Also, `umask 077` is inconvenient to use since it is a global setting per user. For example, if working on a shared project with this setting, the user will have to specifically correct permissions every time a new file is created (a better mechanism would be using a user and directory specific `umask`). Using *richer access control semantics* of having separate permissions for listing directories, creating/removing files only helps in improving the granularity of protection. Without a concrete mechanism of directly and selectively sharing deep rooted data or protecting application data, this mode fails on the rest of the principles. None of the existing *nix tools provide any auditing facilities for data privacy, thus failing on the fifth design principle.

Note that even with a combination of all four mechanisms, there is still inadequate support for privacy protection.

3.2 Comparison with windows access control

While we focus primarily on the *nix access control model, it is interesting to briefly analyse the windows access control model. Windows is most commonly used in a Personal Computer (PC) environment, so data sharing for users on the same system is not very common. However, the Windows NT line of operating systems (NT, Windows 2000, XP) does have multi-user support.

Windows access control is richer than *nix and includes additional permissions for listing contents of a directory (distinct from the *read* permissions), creating files in a directory (distinct from the *write* permission) and more. Also, permissions can be granted at the granularity of individual users. A major difference for our analysis is the fact that in Windows NT/2000/XP Pro, it is possible to share a directory with other users without giving permissions for the complete path to that directory. For example, to share *grand-child* with the path `dir/child/grand-child`, there is no requirement for execute permissions on `dir` or `child`. However, in Windows XP operating system, the *recommended* way (for *XP Home*, it is the only way) of sharing a folder with other users on the same computer is to *move* it to the '*Shared Documents*' folder. This is an undesirable property since it forces the owner to change the directory structure.

For applications, Windows recommends creating private files in the hidden '*Documents and Settings/userName/Application Data*' directory, thus creating a single location which can be more easily protected. However, there is no certain way to ensure that applications follow this principle.

Overall, this access control (except in *XP Home*) provides good support for preventing inadvertent exposure, only partial support for protecting application data and improves the granularity of protection. In regards to convenience, Windows allows a deep rooted directory to be shared easily (no need to set parent permissions). However, it does not let owners represent their data differently for sharing it with other users.

As an additional contrast, we use two design goals in our privacy enhancements. First, we attempt to maintain similar semantics to the basic UNIX model, thus requiring minimal re-education and greater ease of use than the fine grained Windows model which is complex to use. Secondly, our approach gives users the option of selectively using our privacy enhancements with the ability to switching off our proposed enhancements. Thus, it gives a way to slowly introduce VBAC to the user base.

4 Privacy enhancements

The analysis presented in the previous section indicates a need for better access control and monitoring mechanisms. In this section, we present our two solutions that can be used independently or together to facilitate stronger privacy protection in *nix systems. The first solution is a Privacy Auditing Tool that monitors the privacy health of an organisation and can alert users/administrators of potential threats (the fifth principle). The second solution is a new access control model, VBAC, that modifies the data sharing mechanisms of *nix systems and succeeds on the remaining private-sharing principles. Using the two solutions together provides an excellent data sharing environment.

4.1 Privacy auditing tool

The aim of the privacy auditing tool is to periodically monitor user home directories and identify potential private data exposures. A similar approach is used by most enterprise security applications like (Symantec Security Manager, 2008) that audit user systems and enforce compliance to security policies, for example, requiring laptop owners to keep a boot-up password, or system administrators to enforce stricter password rules and so on. In a similar vein, the privacy auditing tool will scan user home directories and alert administrator or the users directly if their private data can be accessed by unauthorised users.

The design of such a tool is very similar to the design of the attack described in Section 2.7. A number of test accounts are created on the monitored system with different group memberships, since it is possible that some user group might have access to more private data than others. The tool is then run from these test accounts to identify exposed private data. The auditing tool can be used either to obtain only higher level statistics, as described in the attack or more user-specific information which can alert users directly of *their private* directories that can be accessed by unauthorised users. A variant would be allowing users to themselves invoke an audit of their home directory. Yet another variant would be to allow the tool to automatically correct some of the obvious mis-configurations like e-mails.

Even though this solution does not solve the underlying access control problem, it has the following *advantages*:

- The privacy auditing tool does not require operating system or file system changes and so can be easily incorporated into enterprise infrastructures. The auditing solution is the quickest way of mitigating this vulnerability.
- Since existing security auditing tools operate in a similar mode, this tool can be easily added on as a privacy protection module to such tools.
- Even with a better access control model, the unavoidable and error-prone human involvement in protecting private data makes such a tool an important component of a secure enterprise.

Next, we discuss a more proactive solution to address the *nix privacy shortcomings.

4.2 View-based access control

Our second solution is the design of a new access control mechanism called the VBAC. Similar to the *nix access control, VBAC is also a discretionary access control model, thus keeping security within the control of the data owner. VBAC is motivated by the private-sharing principles described in Section 3 and is based on the following design goals:

- Act only on primary objects

Private sharing in *nix in its current form requires a two step approach of:

- 1 sharing desired data
- 2 protecting other unrelated data.

The second step, adequately protecting sibling directories or newly created directories, is unintuitive and should be removed. This design feature will help us adhere to the first principle.

- Keep application data only in the owner's view

A severe privacy breach occurred due to improper handling of application profile data. Such data should be viewed only by the owner and unless specifically allowed, should not be visible to other users. This design feature adheres to the second principle.

- Allow hiding of sibling directories from other users

The POSIX ACLs increase granularity of protection for users, allowing data to be shared with individual users. Combining this with an approach that can completely hide sibling data from other users, thus protecting file metadata will comply with the third principle.

- Allow extracting deep rooted directories

In order to share deep rooted directories, it should be possible to simply pluck them from the file system tree and put them in the view of desired users. Also, it should be possible to share a different representation of the data without impacting the owners' view of the file system. This will provide the convenience desired by the fourth principle.

Based on these design goals, the VBAC access control model creates a new file system primitive called a '*view*'. Informally speaking, a data owner can define a view of her home directory, dictating what another user gets to see when he attempts to access it. By adding only the data she wants to share into this view, other data remains protected. Also, it is possible to add a deep rooted directory directly to this view and it can be represented differently. Using such a mechanism, unless the owner explicitly adds her application data into a shared view, it will be always hidden from other users. More details follow in the next section.

5 VBAC: design and implementation

VBAC extends the *nix access control model by adding a *view* primitive, that presents a different file system structure to different users. For every user, there is one owner-view of the home directory which is the same as the standard home directory in current systems. In addition, the owner can define new views of the home directory for other individual users, user groups or the *others* set. For example, a user *bob* can create a view of his home directory for a user *alice* and another view for his user group *faculty* and yet another view for all *other* users. He can then add desired data to appropriate views depending on what he wants to share. The added data could be a deep rooted directory and can be shared using a different name. Other users can access their view of *bob's* home directory using the same *~bob*. The underlying system *automatically routes* them to their appropriate view and users continue to see the view directory as *~bob*. This ensures that no current scripts or access habits are disturbed.

The VBAC model uses the same permission types as baseline *nix – *read*, *write* and *execute* and they have the same semantics. VBAC only adds another layer of access control by making a higher level decision of what a user gets to see or not. After that decision, whatever a particular user has in his/her view, it is access controlled using the baseline *nix permissions. We believe that this feature makes VBAC an elegant extension of the *nix model.

Also, a user can decide to *switch off* the additional VBAC layer providing other users with the same view as the owner-view (of course, access to data is controlled by the lower layer of *nix access control). This implies that VBAC can be incrementally introduced into a system without forcing all users to migrate to it immediately.

Separating the owner view from the views of other users offers us great advantages:

- No sibling directories are put at risk, since only the directories added to user views (that is, the directories that were to be shared) are visible to other users (the first principle).
- No user application files are visible to other users. Only the owner-view contains files like *.history*, *.mozilla*, which were responsible for significant privacy breaches (the second principle).
- File metadata can be protected simply by not adding them to another user's view. If a file is solely in the owner's view, other users cannot access metadata or data or even find out existence of a file. In addition, VBAC uses POSIX ACLs for fine-grained user access control (the third principle).
- As mentioned above, VBAC allows sharing a deep rooted directory without giving permissions throughout the entire path. Also the shared directory can be represented differently in different views (the fourth principle).

To avoid managing views at the granularity of individual users, we provide two mechanisms for doing selective sharing of data by using the group views.

This first mechanism uses Security-by-Obfuscation (SBO). Recall that we mentioned that in order to do selective sharing under an execute-only directory, a simple solution was to use tough-to-guess directory names. However, it was deemed infeasible since the owner would have to keep tens or hundreds of cryptic names. With the separation of views, an owner can now share a particular directory with a tough-to-guess name in the

view, while keeping the original name in the home directory (owner-view). To achieve this, the other users' view is set to execute-only permissions and while adding a directory to the view, the owner also gives a *passphrase* which is used to encrypt the name of the directory being added. The encrypted cipher text is appended to the directory name and the resulting string is used as the target name in the view. Now, the owner authorises a user to access this directory by giving that user the original directory name and the *passphrase*. Notice the similarity with the authorisation mechanism in original execute-only *nix – only one extra piece of information, the passphrase, is delivered using the same out-of-band channels like e-mail. By using cryptographically tough-to-guess names, we are able to bring this user authorisation definition much closer to the underlying system's definition, a shortcoming in the original *nix model.

In order to prevent users having to remember the passphrase every time they need to access another user's shared directory, we provide another mechanism to do selective data sharing using a group view. This second method uses *POSIX ACLs*. In this case, when an owner adds a directory to a view, it is shared with the desired name, but with no permissions for any users other than the owner. Along with that, the directory is flagged and the passphrase is stored in a secure location only accessible by the superuser (it is stored as a trusted extended attribute (Linux Manual pages, 2008)). The owner still authorises the users in the same manner – by informing them of the directory name and the passphrase. However, before accessing the shared directory users have to perform an additional step. They first obtain access by executing a new trusted command (see Section 5.5) that, provided the passphrase was correct, updates the ACL associated with that directory allowing access to that user. All accesses after this step can proceed without having to enter the passphrase. *An added advantage of this method is that all users that access the directory have their names included in the ACL, which can be viewed by the owner, thus serving as an excellent auditing tool.*

5.1 Implementation

In this section, we describe our proof-of-concept implementation of the VBAC access control model.

There are different options using which new access control techniques can be introduced into a system. For example, using a per-process filesystem namespace included in Linux kernels from 2.4.19+ (similar in design to man jails (Linux Manual pages, 2008)) or introducing access control through a Linux Security Module (Wright *et al.*, 2002). We opted for a localised approach that causes minimal impact to the rest of the system. We implemented VBAC through a filesystem based design that allows users to selectively use VBAC on pieces of data that they are most interested to protect. This has an added advantage of allowing to evaluate VBAC independently of the overheads of other security mechanisms.

This new file system is called viewfs and is based on the linux ext2 filesystem (Card *et al.*, 1995). Most of ext2 functions like disk placement of data blocks are reused, viewfs is developed as a loadable kernel module and can be loaded into the kernel without kernel recompilation. viewfs was implemented on a Linux 2.6 kernel. Next, we explain the implementation of important VBAC features. Please note that for some of the details, little familiarity with the Linux Virtual File System (VFS) is helpful. Linux VFS is a file system interface that sits over all underlying filesystems. It provides a default

implementation of most filesystem functions like `lookup`, `mkdir` and also provides an implementation of *nix access control checks. Individual filesystems can choose to override these functions by implementing them in the filesystem.

5.2 View

The foremost VBAC concept is that of a view. From an implementation perspective, a view is a regular directory within the directory containing the owner's home directory (`like/home`). In other words it is a sibling directory to the owner home directory.

However, the view directory has a special name of the form: `'owner.uview.username'` or `'owner.gview.groupname'` or `'owner.oview'`. This name is restricted, that is, users cannot use this name for naming other directories. This restriction helps to identify a directory being a view of another directory and is used to do automatic routing of users to their views. The restriction is enforced in the file system `mkdir` function implementation. The first type of name `'owner.uview.username'` is used to create a view for an individual user. For example, if *bob* creates a view for *alice*, the view will be called `'bob.uview.alice'`. The second type is for a user group and the third is for *other* users. Access is controlled to these views, for example, to prevent user *cathy* from accessing a view for *alice*, by ACLs set at view creation. The period ('.') before the view names keeps them hidden from plain view. It is important to note that the view directory is the parent directory which will contain the data to be shared with other users. As described in Section 5, shared data can be added to this view directory using the security-by-obfuscation or the POSIX ACLs based method.

While it would appear that a more elegant way is to create the view directories within the owner's home directory (no clutter in `/home`), it is not feasible for two reasons:

- 1 First, since we use the same *nix permission types `rwx`, in order to provide access to a view directory within owner's home directory, we would be required to give execute permissions on the home directory, which was the original cause of the privacy breach!
- 2 Secondly, it is easier to implement the view directory as a sibling to home directory due to the VFS implementation of path lookups. VFS looks up a path name `dir/child/grand-child` in a loop by first looking up `child` in the directory `dir` and then looking up `grand-child` in `child`. For routing *alice* to her view of *bob*'s directory, the lookup seeking `bob` in the directory `/home` just returns the entry for `.bob.uview.alice` instead.

Next, we concretely describe the automatic routing to views. A VFS directory lookup occurs in the following manner. Given a name to look up, the directory entry (or *dentry*) cache (or *dcache*) is looked up for the desired directory name. The *dcache* indexes cached *dentries* by hashes of their names. If there is a cache miss, the call passes onto the underlying filesystem for looking up that name. That is followed by the inode number lookup to find the inode number corresponding to the name and if it exists, the inode lookup that gets the object metadata.

We modify this filesystem lookup procedure by first checking if there exists an appropriate view directory. For example, if *alice* is looking for a directory `bob`, we check for the existence of a directory called `'bob.uview.alice'`. At this stage, we have the

complete state necessary for completing this lookup (since view is the sibling to the home directory). If the view exists then the dentry associated with the view directory is returned and is cached in the *dentry* cache on return. In order to work with different user views in the cache, we modify the hash of the dentry by hashing the view name as opposed to the original directory name. This is feasible since the VFS hash function can be overridden by the underlying filesystem.

An important point to note is that this implementation does not interfere with the I/O paths at all, that is when file data is being read or written back to disk. One potential performance impact of the implementation is that a single directory name lookup can cause multiple view name lookups. However, as our initial experiments with benchmarks show, the total overheads are still minimal. Secondly, we foresee *nix installation using viewfs only for user home directories. Therefore the vast majority of system lookups that are to standard OS and other infrastructure files contained in */usr, /etc, /bin* are not affected at all.

As mentioned before, an individual owner can choose to switch off the VBAC model for users accessing her home directory. This is accomplished by keeping an extended attribute (Linux Manual pages, 2008) with the user home directory. The lookup described above first checks for this extended attribute and in case the user chooses to switch off VBAC, the normal ext2 file system lookup is performed providing the basic *nix model.

5.3 Sharing data

The next important viewfs implementation is its mechanisms for adding data to views. It allows adding deep rooted directories directly to user views and possibly with a different name. Also, changes made to the directory in one view should be immediately reflected in all other views. The first idea that comes to mind to facilitate this is directory hard links. A directory hard link is the same inode as the original directory but can have a different name and can be created at any location within the filesystem without worrying about the access along the path to the original directory (unlike a symbolic link). In fact there is no way to distinguish a hard link from the original directory.

However, directory hard links are not allowed by most operating systems including Linux even though it is not mandated by the POSIX standard. The reason is that many OS mechanisms like reference counting and locking consider the file system to be an acyclical tree and hard links can cause cycles. For example, for *a/b/c*, commands *'link e a/b'* and *'link e/c a'* cause a cycle. This will also break many existing applications that traverse the file system assuming it to be a tree. Symbolic links can also cause a cycle but they are easily identifiable since the link is a different inode that stores the complete path to the original pointed-to location. Hard links, as mentioned before, are unidentifiable and cycle detection for hard links can be expensive.

In the context of viewfs however, we can very easily prevent any cycle formation. We can do this by only allowing users to create a link from inside the view pointing outside and never in the other direction. This can be checked while adding data to a view (creating the hard link) and since views have restricted name, such a check would not be expensive. Even though we had a working implementation of this approach, there is an additional issue specific to the linux operating system and its implementation of the VFS dentry cache. This view implementation can break in certain situations, for example,

when a directory and its hard link are both cached and one is deleted. Interested readers can follow Linus Torvald's explanation debating Reiser4 filesystem's directory aliasing at <http://kerneltrap.org/node/3749> (Reiser4 Aliasing Debate, 2008). It is unclear if the same holds for other variants like FreeBSD. However, we preferred linux and opted to take an alternate path.

Similar objectives can be achieved in linux using a bind mount (Linux Manual pages, 2008). A bind-mount mounts one portion of the filesystem tree at another location. Since it is a separate file system mount, the linux implementation issue discussed earlier does not apply (Reiser4 Aliasing Debate, 2008). This provides both sharing the deep rooted directory and sharing it with a different name. For viewfs, an additional requirement is to first create the mount point which will be the desired directory name in the view and is done while adding data to the view. Section 5.5 describes new commands created for facilitating this viewfs implementation and an example viewfs session demonstrating its usage.

5.4 VBAC usability

The following characteristics make VBAC and the viewfs implementation superior in design in regards to usability and integration with existing *nix systems:

- VBAC elegantly complements the *nix access control by only adding a higher level *can-see/cannot-see* decision. Also since the view concept is regularly practiced as an access control tool in databases (Sheth and Larson, 1990; Wang and Spooner, 1987). We believe it will not be tough for users to transition to it.
- Individual users can choose to switch on/off the VBAC access control model, thus allowing incremental introduction of VBAC. In addition, only directories mounted on viewfs are impacted, thus isolating its influence.
- In viewfs, users only need to perform operations on data which they want to share. All other data, including application data, is automatically protected without any explicit user commands.
- VBAC requires only data owners to perform explicit operations. Users accessing shared data continue to do so in normal *nix fashion (like `cd ~bob`), except when using the SBO or ACL private-sharing methods.

5.5 Viewfs example and commands

This section provides a complete list of new viewfs commands and an example usage of the viewfs to share and access data. Following new commands have been implemented for viewfs:

- `createview -[u | g| o] [user | group]`

This commands creates a user/group/others view for the invoker's home directory. It also sets appropriate ACLs to control access to the view. For example, bob executing `'createview -u alice'` or `'createview -o'`.

- `add2view [-u | g] o [user | group] [-b | -a] [src] [target]`
 This command adds the `src` directory to the appropriate view and sets its name to be *target*. Also if `-b` option is used, a passphrase is obtained from the user and the directory is shared using the SBO method described in Section 5. If `-a` option is used, the ACL method is used. For example adding *jobsearch* directory to an *others* view as *CV* using SBO – ‘`addview -o -b jobsearch CV`’. This will prompt the user to enter a passphrase (*pp*) and the directory is shared with the name ‘`CV-{CV}pp`’ where `{}` indicates encryption. To allow ordinary users to perform a mount, this script is setuid root, though root privileges are dropped at initialisation and gained only for the mount operation using `seteuid` (Linux Manual pages, 2008).
- `getacc [target]`
 This setuid command prompts for a passphrase that protects access to the target directory, when using the ACL method. If the passphrase is correct, the ACL of target is modified to give access to the invoking user.
- `vcd [target]`
 This command is to `cd` into a directory protected by the SBO method. It prompts for a passphrase and encrypts the target name with it to obtain the actual directory name. There are also other functions to remove a view, unshare a directory from a view, *etc.*

Figure 2 contains an example `viewfs` session, in which *bob* has created a view for *others* and *alice* is accessing that view.

Lines 2–8 show the contents of *bob*’s home directory. In lines 9, 10 *bob* adds his *personal* directory to *others* view using the ACL method. In lines 11, 12 he adds his *CS6210* directory using the SBO method. Finally he adds his research and web directory in lines 13, 15. Lines 18–25 show the contents of the view directory. Notice the name of the shared *CS6210* directory, the permissions of the shared *personal* directory, as done in the ACL method (see Section 5) and execute-only permissions of the view directory (line 20) for SBO protection.

Alice accesses *bob*’s home directory `~bob`, but cannot list its contents (line 4). She can access the *public_html* and *research* directories, since they have not been protected by SBO or ACL methods. In order to access the *personal* directory, she executes the `getacc` command (line 11) and enters the passphrase, which *bob* would have given her out-of-band. After the command she can access the directory. And listing the ACL of *personal* directory shows the modification allowing access to *alice*. Finally, in lines 25–28, she accesses the *CS6210* directory, protected by the SBO method, using `vcd`.

Figure 2 Example viewfs session

#	BOB	ALICE
0	[bob@local bob] \$ pwd	[alice@local alice] \$ cd ~bob
1	/mnt/vw/homes/bob	[alice@local bob] \$ pwd
2	[bob@localbob] \$ ls -l	/mnt/vw/homes/bob
3	total 7	[alice@local bob] \$ ls -l
4	drwxr-xr-x 2 bob bob 1024 Apr 23 01:19 courses	ls: .: Permission denied
5	drwxr-xr-x 2 bob bob 1024 Mar 26 05:14 mail	[alice@local bob] \$ ls public_html
6	drwxr-xr-x 2 bob bob 1024 Apr 23 20:59 personal	index.html
7	drwxr-xr-x 4 bob bob 1024 Apr 8 17:06 research	[alice@local bob] \$ ls research
8	drwxr-xr-x 2 bob bob 1024 Mar 26 05:16 www	linux-2.6 paper
9	[bob@local bob] \$ add2view -o -a personal personal	[alice@local bob] \$ ls personal
10	Enter Passphrase: ***** (bob12)	ls: personal: Permission denied
11	[bob@local bob] \$ add2view -o -b courses CS6210	[alice@local bob] \$ getacc personal
12	Enter Passphrase: ***** (group4)	Enter Passphrase: ***** (bob12)
13	[bob@local bob] \$ add2view -o research research	[alice@local bob] \$ ls personal
14		passwords
15	[bob@local bob] \$ add2view -o www public_html	[alice@local bob] \$ getfacl personal
16		# file: personal
17	[bob@local bob] \$ cd ../bob.oview/	# owner: bob
18	[bob@local .bob.oview] \$ ls -al	# group: bob
19	total 10	user::rwx
20	drwx- -x- -x 6 bob bob 1024 Apr 23 20:59 .	user:alice:r-x
21	drwx- -x- -x 8 bob bob 1024 Mar 20 21:56 ..	group:: - - -
22	drwxr-xr-x 2 bob bob 1024 Apr 23 01:19 CS6210 -CSYrzy3dw1uIs	mask::r-x
23	drwx- - - -2 bob bob 1024 Apr 23 20:59 personal	other: - - -
24	drwxr-xr-x 2 bob bob 1024 Mar 26 05:16 public_html	
25	drwxr-xr-x 4 bob bob 1024 Apr 8 17:06 research	[alice@local bob] \$ vcd CS6210
26		Enter Passphrase: ***** (group4)
27		[alice@local CS6210 -CSYrzy3dw1uIs] \$ ls
28		intro.6210

5.6 Performance evaluation

We evaluated our viewfs implementation against three popular filesystem benchmarks and compared it with the baseline ext2 performance, the file system viewfs is based upon. The experiments were conducted on a P4 1.6 GHz Dell Inspiron 8200 with 512 MB RAM running RedHat Linux with kernel 2.6.11.3. All results were averaged over multiple runs. It is important to note that the number of users that access a particular view does not have an additional impact on our performance since each user independently

accesses the view directory. The performance will be impacted only due to additional inode lookups required to identify appropriate view directories. To evaluate this, we used two viewfs scenarios:

- 1 *viewfs-owner*: when an owner is accessing her data
- 2 *viewfs-other*: when a user is accessing the data from an *others* view.

Note that additional view name lookups occur only in the latter scenario, so the former scenario is an indication of any other viewfs overheads, for example, of using bind mounts and would also be an estimate of impact on users with switched off VBAC model.

5.7 Andrew benchmark

Andrew benchmark (Howard *et al.*, 1988) is a popular filesystem benchmark. It emulates a software development workload and has five phases:

- 1 creates subdirectories recursively
- 2 copies a source tree
- 3 examines the status of all the files in the tree without examining their data
- 4 examines every byte of data in all the files
- 5 compiles and links the files.

For our experiments, it compiled an OpenSSH-2 client. In the *viewfs-other* implementation the benchmark was run by an *other* user in the owner's view directory. Figure 3 plots the times (in ms) on log scale for each of the phases.

Figure 3 Andrew benchmark results




Table 8 show the overheads of the viewfs-other implementation over ext2. TOTAL is the cumulative time taken to execute all five phases of the Andrew benchmark. The cumulative of these phases is considered to be a representation of a typical software development workload and the difference of 3% on the cumulative Andrew benchmark performance between viewfs and ext2 indicates that for a software development environment, users will have to pay only 3% additional overheads over ext2 while getting the additional privacy features of viewfs. The overheads in the earlier phases are greater since the missed inode lookups to check for view directories form a measurable portion of the total cost (less than 100 ms). However, it is to be noted that the current viewfs implementation is a proof-of-concept prototype and greater optimisations may be possible in future to reduce these overheads. A special area of focus could be use of a cache that can quickly map the user to the right inode for the requested file. This will reduce the number of missed inode lookups. For the later phases, other costs (read/write) are more dominant and additional inode lookups form a very small overhead. As discussed earlier, viewfs does not incur costs for data I/O (read/write). The next set of experiments using the Bonnie benchmark further demonstrate this fact.


Table 8 Andrew benchmark overheads of viewfs-other over ext2

<i>Phase</i>	<i>Overheads (%)</i>
1	36
2	41
3	11
4	9
5	2
<i>Total</i>	3

5.8 Bonnie

In order to test our thesis that viewfs should not impact I/O performance, we evaluated viewfs against ext2 on the Bonnie Benchmark (2008). Bonnie tests the speed of file I/O using standard C library calls. It does reads and writes of blocks in random or sequential order and also evaluates updates to a file. The tests were run for a 400 MB file. Figure 4 shows the results for the three implementation for various I/O modes. The X-axis lists the I/O modes and the Y-axis plots the speeds for those operations. As can be seen, for all read/write modes, there is practically no difference between ext2 and viewfs. This proves that viewfs does not have I/O overheads.

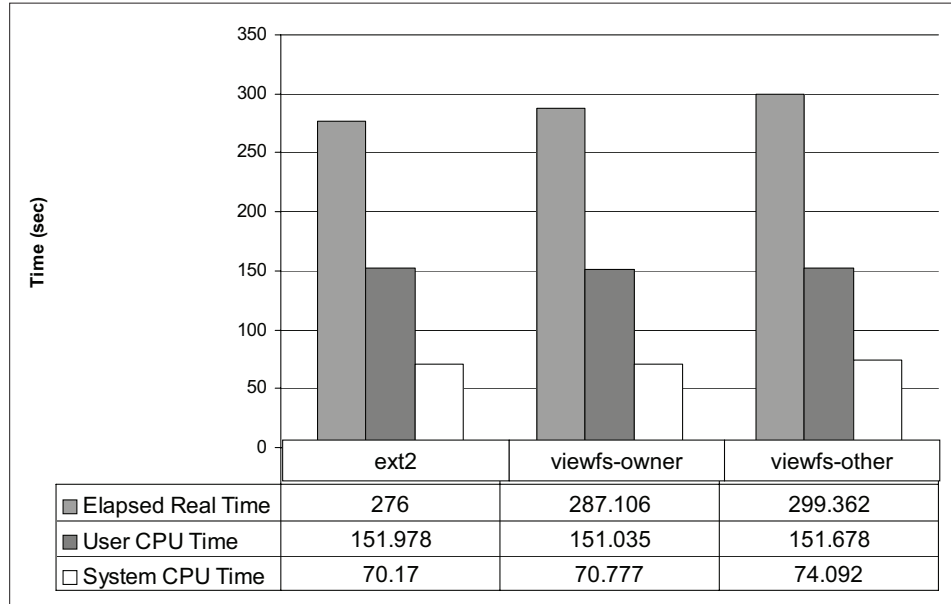
Figure 4 Bonnie benchmark results



5.9 Am-utils

The final benchmark used was the Berkeley Automounter (2008) am-utils. It is a filesystem performance benchmark that configures and compiles the am-utils software package inside a given directory. Overall, this benchmark contains a large number of reads, writes, and file lookups, as well as a fair mix of most other file system operations such as unlink, mkdir, symlink. It is believed to be a good representation of a usual system workload.

Figure 5 shows the results of timing the viewfs-other and viewfs-owner implementations with ext2 using the `time` (Linux Manual pages, 2008) utility. The graph shows that there is only a slight overhead introduced by viewfs (8% in elapsed real time, time between invocation and termination, for viewfs-other). This shows that for typical system usage, viewfs will perform close to the existing filesystems.

Figure 5 Am-utils benchmark results

5.10 Summary of results

These benchmark results show that viewfs performs very well on an average workload. It does have high relative performance overheads for small time operations (for example, metadata intensive operations in Phase 1 and Phase 2 of Andrew benchmark). This is expected since in such cases, the missed inode lookups to check for view directories forms a measurable portion of the total cost.

However, it is important to recognise that these performance costs are paid only for selective pieces of data that require stronger privacy protection. Most storage accessed in systems occur to directories like `/bin`, `/usr/bin` which we neither expect nor recommend to be mounted on viewfs. This is where our data focused perspective on applying privacy enhancements becomes extremely useful. Secondly, these overheads occur only for metadata and viewfs does not impact the I/O paths, as demonstrated using the Bonnie benchmark results above. We firmly believe that these two features make viewfs a reasonable technique from its performance perspective for the enhanced privacy protection that it provides.

6 Related work

We compare our work to three pieces of related work – (1) *nix access control models, (2) alternate technologies like cryptographic file systems, and (3) use of views as access control in other domains.

6.1 *nix access control models

The *nix access control model follows primarily from the original UNIX access control (Ritchie and Thompson, 1974; Ritchie, 1978). It is a Discretionary Access Control (DAC) model and access is granted based on the identity of the subject (user or user group). It is discretionary since access to an object is controlled by the object owners. Most of the research in *nix access control model aims to either improve the granularity of protection (for example, POSIX ACLs (Grunbacher and Nuremberg, 2008)) or counter the threat of malicious programs exploiting the `root` (superuser) privileges. There has been a significant amount of work in developing sandboxing the other similar techniques aimed at allowing only certain filesystem namespace to be visible to an untrusted process (Wagner, 1999; Jaeger and Prakash, 1994; Goldberg *et al.*, 1992; Linux Manual pages, 2008). In contrast, we approach the problem from the underlying *data* perspective. Sandboxing techniques are applied to *processes* irrespective of the data that they are trying to access (even `/usr/bin`, `/bin`, *etc.*). In contrast, the core unit of protection in our approach is *data* and enhancements are applied to only the accesses to such private data. We believe this data perspective to be novel and useful as all processes do not have to comply with privacy enhancements (and pay overheads associated with it).

There also has been significant work at developing new access control models for *nix systems. One variant of Linux developed by National Security Agency (2008), called the Security-Enhanced Linux (SELinux) (2008) supports a mandatory access control model, in which an administrator sets a security policy which is used to determine the access granted to an object and users have limited control on their data. Another access control model is the Role-Based Access Control (RBAC) model (Ferraiolo and Kuhn, 1992; Sandhu *et al.*, 1996), supported by Sun Solaris operating system, in which security attributes can be assigned to user roles (a process or task). This helps reduce the threat of malicious programs exploiting the `root` privileges. In a similar vein, the Rule-Set Based Access Control (RSBAC) model (Ott and Fischer-Hubner, 2001) aims to protect against `root` vulnerabilities and improve granularity of protection. There has also been work on a privacy model (Fischer-Hubner, 1994; 1995; Fischer-Hubner and Ott, 1998) in access control. However, that work is aimed at guiding organisations on how to control their information flow to ensure privacy of collected user-data (for example, healthcare records). To the best of our knowledge this work is the first critical look at the privacy support for *data sharing in a multi-user *nix operating system*. Our proposed VBAC model is a specialised access control model that focuses on *data* and aims at providing stronger privacy protection in such environments. Additionally, we have made a deliberate effort at maintaining the usage of the basic *nix model as much as possible as we believe that it is critical for greater user acceptability.

6.2 Alternate technologies

A common argument against privacy protection mechanisms is that users who care about their privacy would use a stronger security mechanism like encryption. However, sharing encrypted data with many users is a challenging problem. Cryptographic file systems like CFS (Blaze, 1993), CryptFS (Zadok *et al.*, 1998) provide transparent mechanisms of ensuring data confidentiality using cryptographic primitives. Data is stored in an encrypted format on disk and is decrypted (or re-encrypted) on-the-fly while reading from (or writing to) the disk. Another related work is that of Self-Certifying File

Systems (SFS) (Mazières *et al.*, 1999; Mazires, 2000). SFS is used for inter-organisation wide area file sharing aimed at providing a global filesystem image in which users of one organisation can share data with users from another organisation. It uses specialised path names for shared data that allow *self-certification* by deriving them using cryptographic techniques. In practice, cryptographic filesystems have traditionally had poor acceptance in multiuser environments. This is due to the I/O performance effects and also importantly, lack of user education in encryption technologies. In fact a user study reported that even experienced computer users could not use PGP 5.0 in less than 90 minutes and that one-quarter of the test subjects accidentally revealed the secret they were supposed to protect (Whitten and Tygar, 1999). A recent cryptographic file system NCryptFS (Wright *et al.*, 2003) while providing better convenience features, modifies various kernel components like process management, dentry cache and inode cache. This will limit its widespread adoption. Secondly, users have to specifically *attach* to a shared directory, as opposed to continuing to use '~bob' in viewfs. Thus, for its *convenience in use and easier integration with existing systems*, we believe viewfs to be a better suited tool for privacy protection.

6.3 Access control using views

The use of views as an access control tool has been primarily researched in the area of databases (Sheth and Larson, 1990; Wang and Spooner, 1987). Using database *views*, users are only shown the relevant data that they have access to. This is similar in concept to VBAC in which only data that needs to be shared is added to a user's view. A view-based access control model has also been used in networking to control access to management information in the Simple Network Management Protocol (SNMP) (Wijnen *et al.*, 1998). There is also an attempt of creating a new operating system called View-OS (2008) that presents a different view of the system resources including the filesystem, to an OS process. Also, `man` (Linux Manual pages, 2008) can be used to present a different *root (/)* directory to a process. In contrast, our view based mechanism is a comprehensive privacy protection mechanism *preventing many kinds of privacy breaches and works with existing *nix systems*.

7 Conclusions and future work

In this paper, we critically analysed the *nix access control model for privacy support in its data sharing mechanisms. We identified design inadequacies that, combined with user and even application's privacy-indifferent behaviour, lead to privacy breaches. We evaluated two *nix installations of many hundred users and found that a large amount of private data is inadequately protected including e-mails, browsing history data and actual passwords to financial and other sensitive websites. Based on our analysis, we promoted five data sharing principles that should be followed for best privacy protection. Then, we proposed two solutions which when used jointly adhere to all the principles and provide strong privacy protection. As part of the second solution, we propose a new VBAC model which separates the owner's view of the home directory from other users. We also presented a new VBAC-enabled file system, called viewfs. Our experiments with three popular filesystem benchmarks prove that viewfs has acceptable performance, with little overheads.

As part of our future work, we are attempting to deploy viewfs at a real *nix installation and conduct a systematic usability study. We are also investigating the integration of our proposed privacy auditing tool with existing enterprise storage products.

Acknowledgements

This work is partially supported by grants from NSF CyberTrust Program, CSR program, AFOSR program, and IBM SUR grant and an IBM faculty award.

References

- Berkeley Automounter (2008) <http://www.am-utils.org>.
- Blaze, M. (1993) 'A cryptographic file system for Unix', *Proceedings of the ACM Conference on Computer and Communications Security*, pp.9–16.
- Bonnie Benchmark (2008) <http://www.textuality.com/bonnie>.
- Cappelli, D. and Keeney, M. (2008) 'Insider threat: real data on a real problem', CERT Technical report, <http://www.cert.org/nav/present.html>.
- Card, R., Ts'o, T. and Tweedie, S. (1995) 'Design and implementation of the second extended filesystem', *Proceedings of the First Dutch International Symposium on Linux*, ISBN: 90-367-0385-9.
- Carnegie Mellon Software Engineering Institute (2008) <http://www.cert.org>.
- Ferraiolo, D. and Kuhn, D. (1992) 'Role based access control', *Proceedings of the 15th National Computer Security Conference*, pp.554–563.
- Fischer-Hubner, S. (1994) 'Towards a privacy-friendly design and usage of IT-security mechanisms', *Proceedings of the 17th National Computer Security Conference*.
- Fischer-Hubner, S. (1995) 'Considering privacy as a security-aspect: a formal privacy-model', DASY Papers No. 5/95, Computer and System Sciences, Copenhagen Business School.
- Fischer-Hubner, S. and Ott, A. (1998) 'From a formal privacy model to its implementation', *Proceedings of the 21st National Information Systems Security Conference*.
- Goldberg, Y., Safran, M. and Shapiro, E. (1992) 'Active mail: a framework for implementing groupware', *Proceedings of the ACM Conference on Computer-Supported Cooperative Work (CSCW)*, pp.75–83.
- Google Mail (2008) <http://www.gmail.com>.
- Grunbacher, A. and Nuremberg, A. (2008) 'POSIX access control lists on Linux', <http://www.suse.de/~agruen/acl/linux-acls/online>.
- Howard, J.H., Kazar, M.L., Menees, S.G., Nichols, D.A., Satyanarayanan, M., Sidebotham, R.N. and West, M.J. (1988) 'Scale and performance in a distributed file system', *ACM Transactions on Computer Systems*, Vol. 6, No. 1, pp.51–81.
- Jad: JAVa Decompiler (2008) <http://www.kpdus.com/jad.html>.
- Jaeger, T. and Prakash, A. (1994) 'Support for the file system security requirements of computational e-mail systems', *Proceedings of the 2nd ACM Conference on Computer and Communications Security (CCS)*, pp.1–9.
- Linux Manual pages (2008) *Man Command-name*.
- Mazières, D., Kaminsky, M., Kaashoek, M.F. and Witchel, E. (1999) 'Separating key management from file system security', *Proceedings of the ACM Symposium on Operating Systems Principles (SOSP)*, pp.124–139.

- Mazires, D. (2000) 'Self-certifying file system', PhD thesis, MIT.
- Microsoft Hotmail (2008) <http://www.hotmail.com>.
- Mitnick, K., Simon, W. and Wozniak, S. (2002) *The Art of Deception: Controlling the Human Element of Security*, John Wiley and Sons.
- Mozilla (2008) <http://www.mozilla.org>.
- Mozilla Bug Report (2008) 'Bug 59557', https://bugzilla.mozilla.org/show_bug.cgi?id=59557.
- National Security Agency (2008) <http://www.nsa.gov>.
- Opera (2008) <http://www.opera.com>.
- Ott, A. and Fischer-Hubner, S. (2001) *The Rule Set Based Access Control (RSBAC) Framework for Linux*, Karlstad University Studies, Vol. 28.
- Pine (2008) <http://www.Washington.edu/pine/>.
- Reiser4 Aliasing Debate (2008) <http://kerneltrap.org/node/3749>.
- Ritchie, D. (1978) 'The UNIX time-sharing system: a retrospective', *Bell Systems Technical Journal*, Vol. 57, No. 6, pp.1947–1969.
- Ritchie, D. and Thompson, K. (1974) 'The UNIX time-sharing system', *Communications of the ACM*, Vol. 17, No. 7.
- Sandhu, R., Coyne, E., Feinstein, H. and Youman, C. (1996) 'Role-based access control models', *IEEE Computers*, Vol. 29, No. 2, pp.38–47.
- Satyanarayanan, M. (1990) 'Scalable, secure, and highly available distributed file access', *IEEE Computer*, Vol. 23, No. 5.
- Security-Enhanced Linux (2008) <http://www.nsa.gov/selinux>.
- Sheth, A. and Larson, A. (1990) 'Federated database systems for managing distributed, heterogeneous, and autonomous databases', *ACM Computing Surveys*, Vol. 22, No. 3, pp.180–236.
- Sit, E. and Fu, K. (2001) 'Web cookies: not just a privacy risk', *Communications of the ACMU*, Vol. 44, No. 9.
- Slashdot (2008) 'Kevin Mitnick interview', <http://interviews.slashdot.org/article.pl?sid=03/02/04/2233250&mode=nocomment&tid=103&tid=123&tid=172>.
- Symantec Security Manager (2008) <http://enterprisesecurity.symantec.com/products/products.cfm?productid=45>.
- United States Secret Service and CERT Coordination Center (2008) '2004 e-crime watch survey', <http://www.cert.org/nav/allpubs.html>.
- View-OS (2008) 'View-OS: a process with a view', <http://savannah.nongnu.org/projects/view-os>.
- Wagner, D.A. (1999) 'Janus: an approach for confinement of untrusted applications', Technical Report UCB/CSD-99-1056, EECS Department, University of California, Berkeley.
- Wang, C. and Spooner, D. (1987) 'Access control in a heterogeneous distributed database management system', *Proceedings of the Symposium on Reliability in Distributed Software and Database Systems*, pp.84–92.
- Whitten, A. and Tygar, J. (1999) 'Why Johnny can't encrypt: a usability evaluation of PGP 5.0', *Proceedings of USENIX Security*, pp.169–184.
- Wijnen, B., Presuhn, R. and McCloaghrie, K. (1998) 'View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)', *RFC 2275*.
- Wright, C., Cowan, C., Morris, J., Smalley, S. and Kroah-Hartman, G. (2002) 'Linux security modules: general security support for the Linux kernel', *USENIX Security*, pp.17–31.
- Wright, C., Martino, M. and Zadok, E. (2003) 'NCryptfs: a secure and convenient cryptographic file system', *Proceedings of the USENIX Annual Technical Conference*, pp.197–210.
- Zadok, E., Badulescu, I. and Shender, A. (1998) 'Cryptfs: a stackable vnode level encryption file system', Tech Report CUCS-021-98, Computer Science, Columbia University.