Identity Boxing: A New Technique for Consistent System-Wide Identification

Douglas Thain

Technical Report #2005-03
University of Notre Dame
Department of Computer Science and Engineering

Abstract

Today, users of the grid may easily authenticate themselves to computing resources around the world using a public key security infrastructure. However, mapping a user's grid credentials to local accounts has proven to be an administrative hassle. Current techniques for mapping credentials to accounts have a high human burden as well as weak privacy and sharing properties. To remedy this, we introduce the technique of identity boxing. This technique allows a high-level identity to be attached directly to each process and resource that a user employs. Identity boxing eliminates the need to create, manage, and use local accounts: a grid user is known consistently everywhere by his or her credentials. We have implemented identity boxing at the user level within the Parrot interposition agent and applied it to a distributed storage and execution system. The performance overhead of this implementation is only 0.7 to 6.5 percent for a selection of scientific applications, but as high as 35 percent for a metadata-intensive software build.

1 Introduction

Today, the GSI public key security infrastructure allows grid users to be identified with strong cryptographic credentials and and a descriptive, globally-unique name such as /O=UnivNowhere/CN=Fred. [19] This powerful security infrastructure allows users to perform a single login and then access a variety of remote resources on the grid without further authentication steps.

However, once connected to a specific system, a user's grid credentials must somehow be mapped to a local namespace. There are a variety of techniques for performing this mapping. Systems today employ untrusted accounts, private accounts, group accounts, anonymous accounts, and account pools. Each of these methods presents some administrative difficulties. Most techniques must run as the superuser in order to create a new protection domain for the calling user. Many require some explicit interaction with a human administrator in order to generate a new account and update a mapping table. Most permit little or no sharing of data or resources between users on a given system. Large systems such as Grid3 [14] have worked around these problems by employing the old insecure standby of shared user accounts.

Even worse, user identities are not employed consistently across the grid. A single user may be known by a different account name at every single site that he or she accesses, in addition to a variety of identity names given by certificate authorities. In order to access a resource, the user may need to have a local account generated. In order to share resources, each user must know the local identities of users that he/she wishes to share with. However, local identities are often inconsistent or transient, thus preventing any sort of sharing at all.

Ideally, a grid computing system would hide these ugly details from the end user. A user should simply be able to log in and be identified by his or her grid identity without reference to local accounts. If several users wish to share data or resources, they ought to be able to identify each other via their grid identities, rather than by arbitrary local names. This ideal situation is difficult to realize in today's computing systems because of the inflexible nature of the underlying account scheme. Every new user of a grid system must be entered by the administrator into the local account database. Although it is a small burden to do this for one user, it would be a full-time administrative job for systems with hundreds or thousands of users.

To attack these problems, we introduce the technique of identity boxing. This technique is similar to sandboxing: an untrusted program is run by a supervisor that evaluates its actions. The difference is that the identity box attaches a high-level grid identity to every process and resource in the system, without regard to the local account details. This allows a user to execute programs and access data in a coordinated way using only grid identities and ignoring the local account details. Further, the administrator of a resource is relieved of the obligation to create and manage accounts: an identity box can create and destroy protection domains as they are needed. A familiar access control interface allows for the controlled sharing of resources.

We have implemented a prototype identity box using Parrot [42], an interposition agent that provides operating-system-like services at the user level. Because of the user-level implementation, application system calls are penalized by an order of magnitude in latency. This has a marginal overhead on a selection scientific applications, which are slowed down by 0.7 - 6.5 percent in runtime. However, identity boxing is more expensive in meta-data intensive application such as a program build, which is slowed down by 35 percent.

To demonstrate the expressive simplicity of identity boxing, we have employed it within the Chirp [39] personal storage server. The combination of identity boxing with familiar access controls creates a system in which a wide community of users can share resources with little or no intervention by a human administrator.

2 Current Solutions

Figure 1 summarizes methods currently used for mapping grid credentials to local accounts currently in use. Each system has various strengths and weaknesses that we define as follows. A method *requires privilege* if the operator of the service must be the root to employ it. It *protects the owner* if it prevents grid users from harming the service owner after they are admitted. It *allows privacy* if grid users are able to easily protect their data from other users at the same site. It *allows sharing* if grid users are able to easily share their data with others at the same site. It *allows return* if a grid user may store some data, log out, and then log in again at a later time and still be able to access that data. Finally, the *administrative burden* describes how often a human administrator must perform some manual activity to admit grid users.

Single Account. The simplest method of identity mapping is to run all visiting processes in the same account. This method is easy to implement and is often a necessity because it requires no special privileges. Obviously, it does not protect the account holder from malicious users, nor does it afford visiting users any privacy from each other. However, it does allow all users admitted to the account to share data and communicate with each other, if they can be trusted to do so. This approach can be acceptable if it is expected that grid credentials will always correspond to one controlling user. For example, one might reasonably operate a personal GASS file server [7] using only a single account.

Untrusted Account. If it is desired to protect the resource owner from malicious users, a slight variation is to run all processes in a special account for unknown or untrusted users (nobody) that carries fewer privileges than an ordinary user. This approach is generally used by Web and FTP servers. The untrusted account has the same sharing properties as the single account approach, but requires privileges in order to create and use it.

Private Accounts. In systems with distinct users that wish to be protected from one another, one may create a distinct local account for every single user. A table called a "gridmap" file is then needed to map from grid identities to local accounts. This approach was first demonstrated by I-WAY [12] and is widely used today. This approach allows each account to maintain privacy, but does not allow for sharing between accounts. Most importantly, it requires privileges to execute and requires a human administrator to be involved for each new local account creation. In this configuration, the grid credentials are used for securing the connection, but every user still bears the burden of establishing an identity at every site.

Group Accounts. Because of the high administrative burden of creating and maintaining private accounts at every grid site, some systems have turned to creating shared group accounts at every site. This approach is used by the Grid3 [14] system. In this model, there are a small number of accounts, each corresponding to a well-known experiment or collaboration. The involvement of the system administrator is necessary to create the accounts, but once established, multiple users are mapped onto those accounts. These accounts essentially enforce static privacy and sharing policies. Within one group, nothing is private, and all data is shared. Between groups, there is privacy but no

Account Type	Required Privilege	Protect Owner?	Allow Privacy?	Allow Sharing?	Allow Return?	Admin Burden	Example Systems
Single	-	no	no	yes	yes	-	Personal GASS [7]
Untrusted	root	yes	no	yes	yes	per user	WWW, FTP
Private	root	yes	yes	no	yes	per user	I-WAY [12]
Group	root	yes	fixed	fixed	yes	per group	Grid3 [14]
Anon.	root	yes	yes	no	no	-	Condor on NT [43]
Pool	root	yes	yes	no	no	per pool	Globus [18] Legion [26]
Identity Box	-	yes	yes	yes	yes	-	Parrot [42]

Figure 1. Identity Mapping Methods

sharing. As with the other approaches, privileges are required to access multiple accounts.

Anonymous Accounts. As an alternative to group accounts, a system may create a temporary account that lasts only for the duration of a single job. As with private accounts, this requires special privileges, provides privacy, but does not permit sharing. However, it does not require the administrator's involvement for every user. Condor [43] uses this approach on Windows NT by taking advantage of the large numeric user ID space to create a fresh user for every single new job. The primary drawback to this method is that an ID no longer has any meaning after a job completes. Thus, this technique is not suitable for any situation where a job creates persistent data and then must return to it later.

Account Pools. A variation on anonymous accounts may be employed on Unix-like systems. The system administrator may create a pool of anonymous accounts (i.e. `grid0-grid99`) for use by a grid system, allowing a resource manager to assign available accounts to jobs on the fly. This approach is available in both Globus [18] and Legion [26]. Like anonymous accounts, an account pool does not allow for return: a given user might be `grid9` today and `grid33` tomorrow. However, it does protect the system owner from users and users from each other.

Identity Boxing. Identity boxing, as we will explain shortly, dispenses with all of the difficulties of account management that we have described. It allows named protection domains to be created on the fly without reference to any account database. Identity boxing can be employed by any user without root privileges. This allows ordinary users to create grid services without creating new security risks by becoming root. Because each visiting user runs in a protection domain, identity boxing protects the owner from grid users, protects grid users from each other, and allows for both sharing of data, and return to stored data. No administrator intervention is needed to create an identity box.

3 Identity Boxing

An identity box is a well-defined execution space in which all processes and resources are associated with an external identity that need not have any relationship to the set of local accounts. That is, within an identity box, a program runs with an explicit high-level name such as /O=UnivNowhere/CN=Fred rather than with a simple integer UID or account name.

Identity boxing makes it possible to use identities consistently throughout a grid computing system. Regardless of the machine, account, or resources in use, a program and all of its data components use and percieve the same identity everywhere. Permission checks and access control lists are based upon the high-level name rather than low-level account information. Of course, it is necessary to employ *some* low-level account to run programs and store data, but this account name is irrelevant to access control. Further, identity boxing dramatically reduces the administrative burden of operating a grid computing system. Identity boxes can be created at runtime by unprivileged users without consulting or modifying local account databases. A service process within a grid can run without root privileges, eliminating a common source of security problems.

Ideally, identity boxing would be implemented within the operating system kernel. However, as many have observed, practical grid computing requires that we live with unmodified operating systems. Thus, we have implemented identity boxing using an *interposition agent* [29] that provides operating-system-like behavior at the user level without special privileges.

We have modified the Parrot [42] interposition agent to perform identity boxing on arbitrary processes. Parrot works by trapping and modifying the system calls of a child process as it runs. This has been used previously to attach new I/O services to existing applications. For example, Parrot allows a program to access GSI-FTP [2] sites by simply

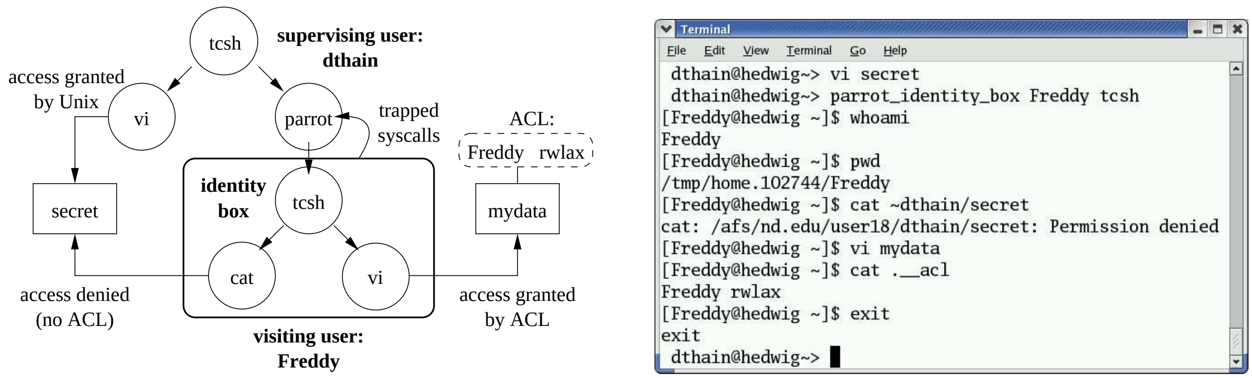

Figure 2. Example of Identity Boxing in an Interactive Session

*An example of identity boxing shown as a schematic and as a shell transcript. The supervising user (*dthain*) creates a file* `secret` *in his home directory. He then creates an identity box for the visiting user* Freddy*, who is not allowed to access* `secret` *because there is no ACL present by default. However,* Freddy *can create a file* `mydata` *in his new home directory, where the ACL has been initialized to give him complete access.*

opening files under the path `/gsiftp`. Internally, Parrot resembles an operating system: it must track a tree of processes, service system calls, and direct I/O requests to drivers. Because of this structure, it is natural to add an operating-system-like feature such as identity boxing. We have modified Parrot to carry with each process a free-form text string indicating the user's high-level identity.

The end-user interface to identity boxing is simple. The user invokes `parrot_identity_box` with an identity string and a command to run. The supervising user can choose absolutely any name for the visitor. MyFriend, JohnQPublic, and Anonymous429 are all valid names. The visiting user retains the supervising user's integer UID, but this is no longer used for access control.

Within an identity box, access control to files and other objects is somewhat complicated because visiting identities are free-form strings. These new identities do not fit into the existing data structures that record UIDs, nor can Parrot modify objects not owned by the supervisor. Our solution to this problem is to abandon the Unix protection scheme and adopt access control lists (ACLs) instead. In each directory, Parrot looks for a file named `.__acl` that describes what actions users can perform on files in that directory. Any program run within an identity box will respect these ACLs. The form of each ACL is similar to those in AFS [24] and similar systems. Each entry lists an identity and the set of operations that can be performed. Identities may contain wildcards in order to match patterns. For example, this ACL allows /O=UnivNowhere/CN=Fred to read, write, list, execute and administer this directory. It also allows any user at /O=UnivNowhere/ to read and list it:

```
/O=UnivNowhere/CN=Fred    rwlax
/O=UnivNowhere/*          rl
```

The visiting user's newly-created home directory is initialized with an appropriate ACL. As in AFS, a newly-created directory inherits the ACL of its parent. Of course, Parrot cannot retroactively place ACLs throughout the file system. When it encounters a directory without an ACL, Parrot enforces Unix permissions as if the visiting user was the Unix user nobody. This ensures that the supervising user's data is protected from the visiting user.

An example of an interactive identity box is shown in Figure 2. Here, the Unix user dthain has created an identity box, naming the contained user Freddy. Note that Freddy does not appear anywhere in the system account list. Freddy attempts to access a file `secret` owned by dthain, but is denied because that file is private to dthain. However, Freddy is given a home directory in which he can work and is allowed to write the file `mydata`.

Figure 2 also shows that the identity box causes the Unix account name to correspond to that of the identity string. This allows `whoami` and similar tools to produce sensible output. This is accomplished by creating a private copy of the `/etc/passwd` file, adding an entry at the top corresponding to the visiting identity, and then redirecting all accesses to `/etc/passwd` to that copy. In addition, a temporary home directory is created for the visiting user's startup files and private data. However, this is merely a convenience. Neither the existing user database nor the private copy play any role in access control within the identity box.

4 Identity Boxing in a Distributed System

Identity boxing has a number of uses in an interactive setting. One might imagine using identity boxes to run untrusted programs downloaded from the web, or perhaps to allow a visitory to briefly borrow one's workstation. However, identity boxing is most useful in a large distributed system where there are many resources and no coordination of accounts and identity. An example of this is the Chirp personal storage server. [39]

A Chirp server is a personal file server for grid computing. It can be deployed by an ordinary user anywhere there is space available in a file system. A Chirp server exports the avaliable file space using a protocol that closely resembles the Unix I/O interface. This file space can be accessed remotely like a distributed filesystem by using Parrot with ordinary applications. A collection of Chirp servers report themselves to a catalog, which then publishes the set of available servers to interested parties.

Of course, there exist a variety of systems for storing data on the grid. GridFTP [2] provides secure, high-performance access to legacy systems. SRB [4] combines databases, file systems, and other archives into a coherent system. SRM [37] defines semantics for storage allocation in time and space. IBP [33] makes storage accessible through a malloc-like interface with access control via capabilities. NeST [6] provides unified access to grid storage through a variety of protocols. However, Chirp is a particularly interesting platform in which to explore identity boxing because it has a fully virtual user space. This means that the space of local users is completely hidden from external users. All data is stored and referenced by external identities.

A Chirp server supports a variety of authentication methods, including Globus GSI [19], Kerberos [38], ordinary Unix names, and a simple hostname scheme. Upon connecting, the client and server negotiate an acceptable authentication method and then the client must prove its identity to the server. If successful, the server then knows the client by a principal name constructed from the authentication method and the proven identity. For example, a single user might be known by any of of these identities:

```
globus:/O=UnivNowhere/CN=Fred
kerberos:fred@nowhere.edu
hostname:laptop.cs.nowhere.edu
```

Once identified, a user may access files on the Chirp server much like any other file server. If the user is employing Parrot, files on a Chirp server appear as ordinary files in the path `/chirp/server/path`. These files are protected by ACLs like those used in Parrot.

Now, imagine the user that wishes to execute a program using data stored on such a server. Traditionally, the user would have to arrange for a login on the same server and use that to access the data directly. However, the user would also have to arrange for the server to store the data under that same identity, which would require the server to run as root. If this was impossible, the user would have to extract the data from the server and run the computation on a different host entirely.

The technique of identity boxing allows to sidestep thse difficulties. To demonstrate this, we have added to the Chirp protocol a simple `exec` call that invokes a remote process. This process is run within an identity box corresponding to the identity negotiated at connection. The identity box enforces access to resources as described above, allowing ordinary applications run unmodified in a remote environment. Of course, the calling user must have the execute (x) right on the program (and any subprograms) to be executed.

The combination of file access control and remote execution allows for simple controls with powerful meaning. If the user has the write and execute (wx) rights on a directory, then he/she can stage in an executable and run it. If the user has only the read and execute (rx) rights on a directory, then he/she is limited to running programs already there. For example, this ACL would allow any user within `nowhere.edu` to run existing programs, while allowing any user holding a `UnivNowhere` certificate to stage in and run any program.

```
/:  hostname:*.nowhere.edu     rlx
    globus:/O=UnivNowhere/*    rwlx
```

This flexibility also introduces new challenges. The ACL given above admits a large set of users that will not necessarily want to share a namespace. Imagine the chaos of allowing one hundred users using the same directory to store files and run programs. Typically, visiting users will require a fresh namespace and the ability to adjust the ACL in order to permit access to their collaborators. For this purpose, an ACL may also include the *reserve right* (V), which is a variation upon amplification. [28] Suppose that the remote users had been given only the reserve right:

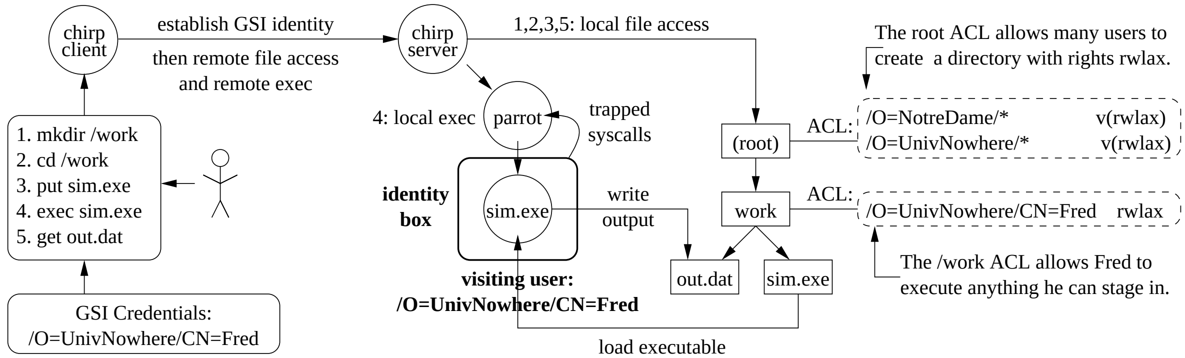

Figure 3. Example of Identity Boxing in a Distributed System

Identity boxing can be used to support visiting users in a distributed system. The Chirp file server provides remote file access and remote file execution to network users. A remote user using a Chirp client creates the `/work` *directory, stages in the* `sim.exe` *program, executes it, and then retrieves the output* `out.dat`*. The Chirp server runs* `sim.exe` *in an identity box corresponding to the remote user. The system may be run by any ordinary user and does not require the creation of any accounts before or during its operation.*

```
/:  hostname:*.nowhere.edu      rlx
    globus:/O=UnivNowhere/*     v(rwlax)
```

When a user performs a `mkdir` in a directory in which he/she only holds the reserve right, the newly-created directory is initialized with an ACL containing the rights listed in parentheses after the V. Not only does this create a private namespace, but it also allows the user to selectively grant access to others. Suppose that the above ACL is present in the root directory when a user identified as `globus:/O=UnivNowhere/CN=Fred` invokes `mkdir(/work)`. The ACL in `/work` would be:

```
/work:  globus:/O=UnivNowhere/Fred    rwlax
```

By virtue of the A right, Fred can further adjust the ACL to give access to other users. Of course, if the system owner does not want a visiting user to extend rights to others, then the A right may simply be left out of the top-level reserve set.

(As an aside, the conscientous reader might wonder how remote users should be expected to choose unique directory names as they create their own workspaces. Although one could imagine server-side support for generating unique directory names, it is sufficient to ask clients to employ the Ethernet approach. [41] Each client may choose a directory name with a random component. In the event of a collision between names, the losing user would not have permission to access the already-existing directory, and should try again after an exponentially increasing delay.)

The combination of identity boxing with a virtual user space and powerful ACLs allows for a dramatically simplified user experience. Given appropriate ACLs, users may discover storage, stage data, run programs, and retrieve output without special privileges or interaction with an administrator. Further, any user is permitted to be a supervisor, deploying and administering any resource that they are able to access. Owners of resources remain in control, delegating and restricting rights as they see fit.

Figure 3 demonstrates how all this fits together. The user `Fred` wishes to run `sim.exe` on a remote machine using his grid credentials. He uses a client tool to contact a Chirp server and creates the `/work` directory using the reserve (V) right. He then stages in the input data and the executable to the remote machine. Using the `exec` call, he invokes the simulation, which is run in an identity box annotated with his name. The identity box allows his simulation to run and access his data securely, even though he does not have an account on the machine. Finally, he retrieves the output and cleans up.

At this point, it is worth pointing out an important aspect of identity boxing. The identity box simplifies the creation and management of protection domains: a system may create an identity box on the fly without regard to any external user database. However, this does *not* mean that identity boxing requires a system to admit arbitrary users. Rather, identity boxing allows a system to have complex admission policies, such as access controls with wildcards, or reference to a community authorization service [32], without the difficulty of reconciling that policy to the existing user database.

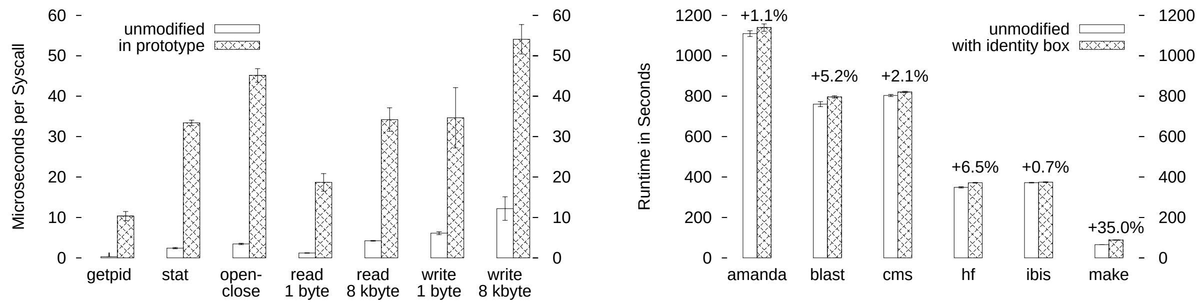

4(a) - System Call Latency

4(b) - Application Runtime

Figure 4. Overhead of Identity Boxing

Within an identity box, individual system calls are slowed by an order of magnitude due to the multiple context switches between the application, the supervisor, and the host kernel. On real applications, the effective overhead varies. A selection of five scientific applications are slowed down from 0.7 to 6.5 percent, but a system-call intensive application such as `make` *is slowed down by 35 percent.*

5 Application Performance

A user-level implementation of identity boxing is fundamentally expensive. In order for Parrot to trap and interpret the system calls of an inferior application, at least six context switches are necessary. Three occur on system call entry: from application to the kernel, kernel to supervisor, and supervisor back to the kernel. Three more are necessary on system call exit. These extra context switches increase latency and also flush processor caches that might otherwise be preserved in an optimized system call mechanism. A more detailed explanation of this overhead is given in an earlier paper. [42]

Figure 4 shows the effects of this performance overhead on individual system calls as well as real applications. Figure 4(a) shows the latency overhead of system calls handled within the identity box. Each entry was measured by a benchmark C program which timed 1000 cycles of 100,000 iterations of various system calls on a 1545 MHz Athlon XP1800 running Linux 2.4.20. Each system call was performed on an existing file in an ext3 filesystem with the file wholly in the system buffer cache. Each call is slowed down by an order of magnitude.

We also run a number of real applications in order to measure the actual overhead of identity boxing amortized over application activity. Five of these were scientific applications that are candidates for execution on grid systems. AMANDA [25] is a simulation of an antarctic gammay-ray telescopic. BLAST [3] searches genomic databases for matching proteins and nucleotides. CMS [23] is a simulation of a high-energy physics apparatus. HF [11] is a simulation of the nucleic and electronic interactions. IBIS [16] is a climate simulation. These applications are described in some detail in an earlier paper. [40] An additional application, make, is simply a build of the Parrot software itself.

The overhead of identity boxing on these applications is shown in Figure 4(b). The five scientific applications are slowed down by only 0.7 - 6.5 percent. Although they are more data intensive than other grid applications, they perform primary large-block I/O. An interactive application such as `make` is slowed down by 35 percent because it make extensive use of small metadata operations such as `stat`. Thus, identity boxing via an interposition agent has overhead that is likely to be acceptable for scientific applications, especially if the technique empowers the user to harness a larger array of resources.

6 Related Work

Sandboxing. Identity boxing is closely related to sandboxing. A sandbox runs an untrusted program underneath a supervisor process which traps its operations and checks them with a reference monitor. The mechanism can be binary rewriting, as in Shepherd [30], a kernel module, as in Janus [22], or the debugging interface, as in Systrace [34]. These systems all require the user must state a list of acceptable operations. Another possibility is to associate rights with programs rather than users, as in SubDomain [10] and MAPBox [1]. Ostia [21] delegates all operations to an agent, allowing for arbitrary policies. One might also consider the Unix `chroot` mechanism to be a simplified sandbox. `chroot` creates a fresh, empty file space in which an application can work but not escape.

Traditional sandboxing requires users to provide some specification or interactive filtering of the system calls attempted by an application. This is an enormous burden because most users have no idea what happens deep within an application. For example, a user running a word processor thinks (quite logically) that the word processor only needs to read and write the file that he/she is editing. In fact, the program actually needs to load an executable, read a configuration file, load plugin libraries, access the dynamic linker, read the host database, create backup files, and use a whole host of other resources that the user has never heard of and doesn't understand. In our field experience with scientific applications [42, 40, 5], even authors of technical software are surprised to learn exactly what system calls their programs attempt. Users are insulated from the system by so many layers of software that we cannot expect them to think in terms of low-level system calls.

Identity boxing builds upon sandboxing by providing built-in access controls that correspond to familiar concepts. Rather than requiring the supervisor to state the access control policy in advance, identity boxing allows the visiting user to interact with others as a first class citizen.

Privilege Separation [35] attacks the same problem in a different way. Many programs, such as login servers, only need some subset of the superuser's capabilities. A common subset is simply the ability to call `setuid()`. However, the sheer complexity of a login server makes it difficult to trust the entire program. Thus, the server itself can be run in an untrusted mode. When it requires a privileged operation, it must explicitly request it from a small kernel of privileged code, which checks the intended operation and then performs it on behalf of the server. This technique is powerful and effective, but still requires a small amount of privileged code and perhaps some code transformation. [8] Identity boxing provides the same power as privilege separation, but requires no privileged code at all.

Virtual Machines. The virtual machine has been proposed as the solution to a variety of problems in distributed computing [36, 44], grid computing [15, 9], operating system composition [17, 27, 13], and security[20, 31]. A virtual machine can completely isolate a service provider from the contained user. This provides both security and an unrestricted workspace for the contained user, who can safely be an administrator in the virtual environment. This is enormously useful ability, particularly when developing a new operating system or performing whole-system simulation.

One could use an entire virtual machine for the same purpose as an identity box. Although this is *possible*, it is quite *impractical* for two reasons. First, creating a virtual machine is a non-trivial administrative activity: one must generate disk images, setup user databases, and install software within the virtual machine itself. Effectively, the creation and management of virtual machines is an activity only accessible to those already skilled in system administration. This also may come at a significant performance cost to move data in and out of the virtual machine. Second, the virtual machine inhibits sharing where it is most needed. Users that run untrusted programs generally *want* those programs to interact with the existing system in a limited way. They want to retain access to local files, to interact with existing processes, to communicate over the existing network. Virtual machines deliberately isolate visiting users, while the identity box encourages controlled sharing.

7 Conclusion

Much of the difficulty of grid computing arises from the problem of mapping our desires onto the limitations of current systems. We wish to employ grid credentials as user identities, but we must also work with existing account names. We wish to employ new forms of storage, but applications still require the ordinary filesystem interface. We wish to execute programs on remote resources, but users desire to work with familiar interfaces.

In this paper, we have shown how the new and the old can work together relatively transparently. By interposing software at both the client side (Parrot) and the server side (Chirp), we are able to preserve existing programs, data, and interaction models while introducing new identities and sharing policies.

References

[1] A. Acharya and M. Raje. MAPbox: Using parameterized behavior classes to confine applications. Technical Report UCSB TRCS99-15, University of California at Santa Barbara, Computer Science Department, 1999.

[2] W. Allcock, A. Chervenak, I. Foster, C. Kesselman, and S. Tuecke. Protocols and services for distributed data-intensive science. In *Proceedings of Advanced Computing and Analysis Techniques in Physics Research*, pages 161–163, 2000.

[3] S. Altschul, W. Gish, W. Miller, E. Myers, and D. Lipman. Basic local alignment search tool. *Journal of Molecular Biology*, 3(215):403–410, Oct 1990.

[4] C. Baru, R. Moore, A. Rajasekar, and M. Wan. The SDSC storage resource broker. In *Proceedings of CASCON*, Toronto, Canada, 1998.

[5] J. Bent, D. Thain, A. C. Arpaci-Dusseau, R. H. Arpaci-Dusseau, and M. Livny. Explicit control in a batch-aware distributed file system. In *USENIX Networked Systems Design and Implementation*, 2004.

[6] J. Bent, V. Venkataramani, N. LeRoy, A. Roy, J. Stanley, A. Arpaci-Dusseau, R. Arpaci-Dusseau, and M. Livny. Flexibility, manageability, and performance in a grid storage appliance. In *Proceedings of the Eleventh IEEE Symposium on High Performance Distributed Computing*, Edinburgh, Scotland, July 2002.

[7] J. Bester, I. Foster, C. Kesselman, J. Tedesco, and S. Tuecke. GASS: A data movement and access service for wide area computing systems. In *6th Workshop on I/O in Parallel and Distributed Systems*, May 1999.

[8] D. Brumley and D. Song. Privtrans: Automatically partitioning programs for privilege separation. In *USENIX Security Symposium*, August 2004.

[9] J. Chase, L. Grit, D. Irwin, J. Moore, and S. Sprenkle. Dynamic virtual clusters in a grid computing environment. In *High Performance Distributed Computing*, June 2003.

[10] C. Cowan, S. Beattie, G. Kroah-Hartman, C. Pu, P. Wagle, and V. Gligor. Subdomain: Parsimonious server security. In *USENIX Systems Administration Conference*, 2000.

[11] P. E. Crandall, R. A. Aydt, A. A. Chien, and D. A. Reed. Input/output characteristics of scalable parallel applications. In *Proceedings of the IEEE/ACM Conference on Supercomputing*, San Diego, California, 1995.

[12] T. A. DeFanti, I. Foster, M. E. Papka, and R. Stevens. Overview of the I-WAY: Wide area visual supercomputing. *International Journal of Supercomputer Applications*, 10(2/3):121–131, 1996.

[13] B. Dragovic, K. Fraser, S. Hand, T. Harris, A. Ho, I. Pratt, A. Warfield, P. Barham, and R. Neugebauer. Xen and the art of virtualization. In *Symposium on Operating Systems Principles*, 2003.

[14] R. G. et al. The Grid2003 production grid: Principles and practice. In *IEEE Symposium on High Performance Distributed Computing*, 2004.

[15] R. J. Figueiredo, P. A. Dinda, and J. A. B. Fortes. A case for grid computing on virtual machines. In *International Conference on Distributed Computing Systems*, May 2003.

[16] J. Foley. An integrated biosphere model of land surface processes, terrestrial carbon balance, and vegetation dynamics. *Global Biogeochemical Cycles*, 10(4):603–628, 1996.

[17] B. Ford, M. Hibler, J. Lepreau, P. Tullmann, G. Back, and S. Clawson. Microkernels meet recursive virtual machines. In *Operating Systems Design and Implementation*, 1996.

[18] I. Foster and C. Kesselman. Globus: A metacomputing intrastructure toolkit. *International Journal of Supercomputer Applications*, 11(2):115–128, 1997.

[19] I. Foster, C. Kesselman, G. Tsudik, and S. Tuecke. A security architecture for computational grids. In *ACM Conference on Computer and Communications Security Conference*, 1998.

[20] T. Garfinkel, B. Pfaff, J. Chow, M. Rosenblum, and D. Boneh. Terra: A virtual machine-based platform for trusted computing. In *Symposium on Operating Systems Principles*, 2003.

[21] T. Garfinkel, B. Pfaff, and M. Rosenblum. Ostia: A delegating architecture for secure system call interposition. In *Symposium on Network and Distributed System Security*, 2004.

[22] I. Goldberg, D. Wagner, R. Thomas, and E. A. Brewer. A secure environment for untrusted helper applications. In *USENIX Security Symposium*, San Jose, CA, 1996.

[23] K. Holtman. CMS data grid system overview and requirements. CMS Note 2001/037, CERN, July 2001.

[24] J. Howard, M. Kazar, S. Menees, D. Nichols, M. Satyanarayanan, R. Sidebotham, and M. West. Scale and performance in a distributed file system. *ACM Transactions on Computer Systems*, 6(1):51–81, February 1988.

[25] P. Hulith. The AMANDA experiment. In *Proceedings of the XVII International Conference on Neutrino Physics and Astrophysics*, Helsinki, Finland, June 1996.

[26] M. Humphrey, F. Knabe, A. Ferrari, and A. Grimshaw. Accountability and control of process creation in metasystems. In *Network and Distributed System Security Symposium*, February 2000.

[27] S. Ioannidis and S. M. Bellovin. Sub-operating systems: A new approach to application security. In *SIGOPS European Workshop*, February 2000.

[28] A. K. Jones and W. A. Wulf. Towards the design of secure systems. *Software - Practice and Experience*, 5(4):321–336, 1975.

[29] M. Jones. Interposition agents: Transparently interposing user code at the system interface. In *Proceedings of the 14th ACM Symposium on Operating Systems Principles*, pages 80–93, 1993.

[30] V. L. Kiriansky. Secure execution environment via program shepherding. In *USENIX Security Symposium*, August 2002.

[31] M. Laureano, C. Maziero, and E. Jamhour. Intrusion detection in virtual machine environments. In *EUROMICRO Conference*, September 2004.

[32] L. Pearlman, V. Welch, I. Foster, C. Kesselman, and S. Tuecke. A community authorization service for group collaboration. In *IEEE Workshop on Policies for Distributed Systems and Networks*, 2002.

[33] J. Plank, M. Beck, W. Elwasif, T. Moore, M. Swany, and R. Wolski. The Internet Backplane Protocol: Storage in the network. In *Proceedings of the Network Storage Symposium*, 1999.

[34] N. Provos. Improving host security with system call policies. In *USENIX Security Symposium*, August 2004.

[35] N. Provos and M. Friedl. Preventing privilege escalation. In *USENIX Security Symposium*, August 2003.

[36] C. P. Sapuntzakis, R. Chandra, B. Pfaff, J. Chow, M. S. Lam, and M. Rosenblum. Optimizing the migration of virtual computers. In *Symposium on Operating Systems Design and Implementation*, 2002.

[37] A. Shoshani, A. Sim, and J. Gu. Storage resource managers: Middleware components for grid storage. In *Proceedings of the Nineteenth IEEE Symposium on Mass Storage Systems*, 2002.

[38] J. Steiner, C. Neuman, and J. I. Schiller. Kerberos: An authentication service for open network systems. In *Proceedings of the USENIX Winter Technical Conference*, pages 191–200, 1988.

[39] D. Thain. Chirp: An architecture for cooperative storage. Technical Report 2005-02, University of Notre Dame, Computer Science and Engineering Department, February 2005.

[40] D. Thain, J. Bent, A. Arpaci-Dusseau, R. Arpaci-Dusseau, and M. Livny. Pipeline and batch sharing in grid workloads. In *Proceedings of the Twelfth IEEE Symposium on High Performance Distributed Computing*, Seattle, WA, June 2003.

[41] D. Thain and M. Livny. The ethernet approach to grid computing. In *Proceedings of the Twelfth IEEE Symposium on High Performance Distributed Computing*, Seattle, WA, June 2003.

[42] D. Thain and M. Livny. Parrot: Transparent user-level middleware for data-intensive computing. In *Proceedings of the Workshop on Adaptive Grid Middleware*, New Orleans, September 2003.

[43] D. Thain, T. Tannenbaum, and M. Livny. Condor and the grid. In F. Berman, G. Fox, and T. Hay, editors, *Grid Computing: Making the Global Infrastructure a Reality*. John Wiley, 2003.

[44] A. Whitaker, M. Shaw, and S. D. Gribble. Denali: Lightweight virtual machines for distributed and networked applications. In *USENIX Annual Technical Conference*, June 2002.