Empirical Privilege Profiling

Carla Marceau and Rob Joyce
ATC-NY

Abstract

The well-known Principle of Least Privilege states that a program should run with the minimal authority that it requires to get the job done, and no more. However, application of the principle has been left to software developers, developers of installation procedures, and system administrators with few tools to assist them. How much privilege does a given program need? How do you know if you write a program that uses too much privilege or install a program with too little? Empirical privilege profiling provides a partial answer to this question by tracking a program's actual use of resources, which can be used as a guide during program development and installation. In this paper, we introduce the concept of dealing with privilege as a measurable quantity, rather than in terms of a "rule of thumb."

1 Introduction

The Principle of Least Privilege [1] states that a program should run with the minimal authority that it requires to get the job done, and no more. Programs, sites, and organizations that observe this principle help to minimize the amount of damage that can be caused by errors in a program or attacks that subvert it. For thirty years, this principle has served as a guide to program and system developers. However, it is difficult to translate this principle into practice. In general, a program will always have more privilege than it needs: if it has less, it will fail, people will be alerted, and the program privilege will be increased (perhaps more than necessary). By contrast, it typically goes unnoticed when a program has too much privilege. Worse, developers often take shortcuts in order to ease implementation, debugging, and testing, sometimes resulting in poor default installation parameters and hidden backdoors in programs. These can be exploited by either insider or external attackers.

Many years ago, one of the authors worked on the Multics project at MIT. As code was developed, team members specified access requirements for modules of the operating system, reasoning carefully about what access was needed by whom. When the system was finally tested, it failed repeatedly, because few modules had sufficient privilege to do their work. In order to get the system running, the programmers changed all access control lists to give sweeping access rights to files—including even "execute" permission on data files! Although this occurred before the formulation of the principle of least privilege—and may have contributed to it, it illustrates the difficulty of getting privilege right without tools.

We have begun an effort to create a system that can empirically approximate the least privilege required by a program. We assume that the program is a black box in the sense that we cannot examine and instrument the source code. (Some recent papers [2, 3] distinguish between white-box, gray-box, and black-box based on whether the analysis delves into program source or

binary, examines artifacts such as the program stack, or can look only at kernel calls. By "black-box," we simply mean that source code is not available and that analysis focuses on execution, not on program artifacts.) The general approach is to run many instances of the program, typically with multiple users, multiple hosts, and multiple sites, and to record the privileges actually exercised by all these instances of the program. This information is then collected to create an abstract composite *privilege profile* for the program. Any single program instance may well use less privilege than the composite profile; however, the composite profile specifies a reasonable minimum, and quite possibly less than granted to the program by the default installation.

Empirical privilege profiling is potentially of use in at least three areas. First, groups of users can profile the privileges that a given application actually exercises when they use it; the profile for a group that uses the application for a limited purpose or in a restricted way could be quite different from the profile of a group of, say, undergraduate students. System administrators could use that information to set up the program's access to resources. Second, program developers could profile the program's use of resources and look for anomalies that indicate poor or excessive use of resources. Finally, tracking resources used by a program could provide a rich new data stream for intrusion detection.

A naïve approach to profiling program privilege would be to list the raw privileges exercised by the program, such as "read file `C:\foo\bar\baz`". However, raw privileges used in specific instances of a program are a poor indicator of future use. It is necessary to generalize over a wide range of individual computers, file systems, and sites; raw file names (and other site-specific privilege information) are not acceptable. To collect *abstract program privileges* it is necessary to correlate privileges exercised by the same program on different computers at different sites and create an abstraction of each privilege.

Further, a naïve approach to profiling program privilege would collect a *set* of privileges exercised by the program. However, a simple set does not provide enough information for correlating privileges across distinct program instances. A program may write to file A and file B, for example, where file A is a file specified by a user, while file B is a file private to the program, whose integrity is crucial to correct operation. Or file A might be a file that is only read by the program, while file B is one that it writes. In correlating privilege exercised by two programs, it is important to correlate the "same" privileges: A_1 with A_2 (instances of A written by programs P_1 and P_2, respectively) and B_1 with B_2 (similar instances of B).

Our approach to privilege correlation is to associate each exercise of privilege with a point in the program. A point in the program that reads a preferences file, for example, will read a preferences file in every instantiation of the program at every site. Therefore, we construct a model of the program's *behavior*; we then associate each exercise of privilege with a program point in the model. The intuition is that the set of program points that access a given resource captures the semantics of the resource *from the point of view of the program*.

To investigate the feasibility of program privilege profiling, we have instrumented Windows programs to collect data about their use of resources, implemented our approach in a simple test bed, and conducted an experiment to create an abstract profile based on several instances of

Microsoft Notepad. In this paper, we report on the results of this preliminary investigation of empirical privilege profiling.

2 A model of program privilege in Windows

A *privilege profile* for a program consists of a set of *privileges* that the program must have in order to function properly. A *privilege* is a *resource* together with a set of access rights on that resource.

We distinguish between *abstract* and *concrete privileges*; more precisely, the distinction is between abstract and concrete *resources*. A concrete resource is a file, registry key, or other resource, that the program accesses by a specific name such as a pathname. An abstract privilege is the use of an abstract resource—that is, an abstraction of concrete resources. Intuitively, an example of an abstract resource might be "the log file for the application" or "any file that the application creates at the request of the user."

In order to create abstractions for concrete privileges used in many instances of an application, we must correlate instances of the "same" concrete privileges. The central hypothesis of this effort is that such correlation is possible using the *program points* at which privilege is exercised. Specifically, we reason as follows:

- A concrete resource that is an instance of a given abstract resource will be accessed at the same program points as other instances of that abstract resource
- Therefore, by comparing the program points at which concrete resources are used, we can abstract from those instances to an abstract resource

Because program points are central to this line of reasoning, a privilege profile also includes the program points at which each (abstract) privilege is exercised.

In the following sections, we further define the concepts of privilege, resource, and program points. We describe how these concepts are implemented in Windows and how we model them. Finally, in Section 2.3, we summarize the contents of the privilege profile of a program.

2.1 Privilege

In this section, we describe first how the Windows 2000/XP operating systems define program privilege and then our model of program privilege.

2.1.1 Privilege in Windows

The Windows security model is more complicated than that of UNIX. Windows distinguishes between access rights, which apply to *securable objects* and are documented in access control lists, and "privileges," such as `SeDebugPrivilege` (ability to modify the memory of another process) or `SeBackupPrivilege` (copy files for backup), which control access to system resources and system-related tasks and are therefore needed to perform *privileged operations*.[1]

[1] In our work, we create a uniform model of privilege by treating such Windows privileges as access rights on a generic SYSTEM resource.

System privileges are assigned by administrators to user and group accounts. In this paper, we will not be further concerned with such special privileges, but in general they must be considered when creating privilege profiles for programs.

Windows defines access rights for many kinds of resources. Persistent resources include files, directories, registry keys, and system services (such as a service that provides access to the Internet). Ephemeral resources include pipes, threads, inter-process synchronization objects, access tokens (which capture salient information about the process's owner), and many others. We are concerned only with privileges on persistent objects.

Each type of resource is associated with a set of type-specific access rights; for example, FILE_READ_ACCESS is an access right on files and directories. Type-specific rights are meaningful to objects of a specific class; for example, KEY_QUERY_VALUE applies to registry keys and FILE_APPEND_DATA to files and directories. Windows also defines generic access rights (GENERIC_READ, GENERIC_WRITE, and GENERIC_EXECUTE), which apply to all classes of objects and are mapped by the kernel into type-specific access rights for each class of securable object. Users can also define new resource types and new types of access rights on them. The work described in this paper focuses on file and directory resources. However, it applies equally to registry keys and we believe can be extended to all persistent resources.

Each kernel operation requires a certain set of rights to succeed. We instrument calls to the kernel at the interface to `ntdll.dll`, which is the Windows interface to the kernel. In some cases, the required set of rights depends on arguments to the operation. For example, the operation to open a file specifies requested rights on the file; if the owner of the process does not enjoy those rights, the operation fails.[2]

For securable objects, we define a privilege as *a set of access rights associated with a securable object*. For example, the right to read file Foo can be represented as `<Foo,FILE_READ_ACCESS>`.[3] Access rights for a given securable object are defined by the object's owner in a Discretionary Access Control List (DACL) associated with the object.

2.1.2 Constant and variable resources

It is important to distinguish between constant and variable abstract resources. A *constant* or *program-specified resource* is one that the program "knows" about, such as a configuration file, a home directory, or a DLL that the program needs to execute. The exact pathname of a configuration file may vary between sites, for example based on the value of an environment variable. Nevertheless, the program typically expects and knows how to find this resource. At one or more points in the program, for example, the program will read, execute, or write a "constant" file resource; when the program accesses the file, it is always from these points. A (resource type-specific) invariant property is associated with a constant resource; for example, every instance of a constant file resource will have the same filename (last component of the

[2] Many `ntdll` calls to open an object also accept MAXIMUM_ALLOWED as a desired access; in such cases, the rights actually required depend on the operations (e.g., NtWriteFile) performed on that object's handle.
[3] In general, the second element of the pair is a set of access rights. Reading a file in Windows typically exercises five distinct access rights.

pathname). By using additional information, such as environment variables, we can specify invariant properties more accurately and completely.

By contrast, a *variable* or *parametric resource* is one that the program has no a priori notion of. For example, the user may provide a file to the program as a parameter. In this case, the pathname and content of the file have no meaning to the program; any file supplied by the user will do. We call this a variable resource, because it typically takes on different values for different executions of the program; it can also take on different values within a single execution. Other variable resources may be determined by the system configuration. For example, a file dialog box may always access the user's home directory.

Distinguishing between constant and variable resources is necessary in order to accurately portray the program's use of resources. The variables are like parameters to a function: a variable can take on many distinct values (concrete resources), even within the same instantiation of the program. If Notepad is used to edit three files, each file corresponds to the same abstract variable resource. Advice to a system administrator about the access required for these resources must be couched in general terms: "Any file edited in Notepad requires …." By contrast, each constant resource can be named more or less accurately: "The initialization file, which is in the application's home directory and is named '`foo.ini`' requires …."

The distinction between constant and variable resources can be exploited for correlating privileges between different executions of a program. A program's initialization file, for example, always has the same filename (local name in the directory) and is always in the "same" place in the file system (where "same" is typically relative to a home directory or other parameter); correlation can exploit these invariant properties. By contrast, variable data files have no *a priori* similarity, other than the way in which the program accesses them.

2.2 Program points

It remains to make precise our intuitive notion of "program point," which should be tied to both observations of program behavior (enabling us to peek inside the "black box") and to the application program to be profiled. We capture these by defining a program point to be a pair consisting of the application program counter and the ntdll operation.[4] Intuitively, the application program counter captures *what the program thinks it is doing with the resource*. The ntdll operation captures its *use* of the resource and the *privileges* required for that use.

Note that because most Windows programs—and indeed, modern programs in general—make extensive use of DLLs, a single application program counter typically causes many ntdll operations to be invoked. We have observed dozens of different ntdll operations associated with a single application program counter.

[4] The typical application program counter is a program counter within the `.exe` file that launches the application. However, applications can include their own DLLs. In this paper, we emphasize the intuitive notion of a control point in the application code.

2.3 The privilege profile

An individual exercise of privilege is a pair <resource, access right>. For example, the right to read file Foo is represented as `<Foo,FILE_READ_ACCESS>`. The heart of the privilege profile is a set of privileges.

In order to facilitate correlation of concrete resources, a privilege profile records privileges together with the program points at which they are exercised. For each resource, the profile lists the program points that access the resource; there may be one or more of these. For each program point, the profile documents the access rights required for the access.

More precisely, let us denote a map (partial function) from range R to domain D by {R => D}. Then the heart of the program privilege profile is a map:

```
{Resource => {ProgramPoint => AccessRight}}
```

(In particular, note the difference between this and two maps

```
{Resource => ProgramPoint}
{ProgramPoint => AccessRight},
```

which would not distinguish between access rights required on two different resources that are accessed at the same program point.)

The profile also partitions resource into three groups:

Constants. These abstract resources correspond to concrete resources accessed by the program for its own use. They include DLL and program files, initialization files, and the like.

Variables. These abstract resources correspond to other files that the program may need to access in order to perform work on behalf of the user. For example, a file that the user wishes to edit is a variable resource.

Uncategorized (Limbo). To avoid errors and malicious inputs, one might require N instances of a resource from M different sites to create an abstraction (see Section 4.2). When the EPP profiler first encounters a concrete resource of a new kind, it does not know how to classify it. Not yet categorized concrete resources are maintained separately, until sufficient information is available that constant or variable abstract resources can represent them.

We can construct a profile by first running the program several times on test data to obtain a preliminary profile. The preliminary profile will contain some constants, some variables, and some concrete resources that have not occurred often enough to be categorized. As more data is collected, the EPP profiler incrementally learns more about these resources—and new ones that occur during the collection—and will be able to abstract them as either constant or variable abstract resources.

3 An experiment in privilege abstraction

We performed an experiment to test our ideas of privilege abstraction. Three different users each exercised Microsoft Notepad, performing various functions for a short time, for a total of eight "runs."

Some parameters of the experiment are shown in Table 1. Notepad's use of resources was captured and a log of exercises of privilege was prepared for each run. The logs were fed into a correlation engine that attempted to correlate the runs' use of resources by comparing program points. We required three different concrete instances of an abstract resource to categorize the resource as a variable or constant. In addition to program points we used limited additional information for correlation; specifically, we identified DLL files to the correlation procedure as constant resources. In a practical system, invariant properties of all constant resources could be extracted from the binary code of the program, the DLLs it invokes, and possibly from other sources, such as environment variables of the platform. The invariant properties would be used to ensure that (a) no variable resource was incorrectly categorized as a constant, and (b) no constant (from that run) was interpreted as a variable.

Run	User	Log size (KB)[5]	# uses of privilege	# resources used	Actions
1	Rob	3588	951	281	Create new file on desktop
2	Rob	4042	1049	64	Create 2 files on desktop
3	Test	1966	518	56	Open (for reading) one file from command line and one from GUI
4	Test	1523	406	51	Open file for reading, save as new file (create file)
5	Test	2148	578	54	Open 2 files for reading and create 2 new files
6	Carla	2476	715	65	Create new file and print it
7	Test	51	20	5	Create and write file from command line
8	Carla	3539	963	57	Read, edit, and save file

Table 1. Use of file resources in eight invocations of Notepad

In total, 633 concrete resources were accessed at thirteen program counter values of Notepad.[6] Correlation identified 53 concrete abstract resources (many of them DLLs) and 6 variable

[5] Log sizes are very large because logs are written in XML.

abstract resources. After all eight instantiations had been analyzed, 44 concrete resources remained unassigned.

Correct assignments. Of the 53 constants identified by correlation, 49 were correctly identified. All six variable resources were correctly identified, representing from three to eight distinct files or directories each.

Variables identified as constants. Four resources that clearly are instances of variables were identified as constants by our algorithm. Two others (.exe files) are probably variables, yet they happened to occur in many of our runs, due to the small number of test sites. Mis-assignments occurred, for example, when common directories (e.g., "`c:\Documents and Settings\All Users`") occurred in almost every run.[7] These mis-assignments could be avoided by identifying the invariant properties of all constants in advance.

Constants identified as variables. None.

Resources that the algorithm was unable to classify. The algorithm was unable to identify 44 concrete resources of the 633 that it encountered, due to insufficient evidence. For example, only one run used the printing function, which invokes DLLs not otherwise used; the algorithm therefore lacked enough evidence to classify them. In practice, we would expect that at the beginning of privilege profiling, there would be many unclassified resources, but that after a while very few would remain.

4 Protecting collaborator privacy and detecting incorrect inputs

We now consider how to protect the privacy of collaborators that provide input to the profile and to protect the profile from malicious inputs.

4.1 Protecting collaborator privacy

The privilege profile contains a list of resources, with the privileges that are required for each resource. The potential threat to privacy is that the names of resources could reveal information about the collaborator's site that is not intended to be public. For example, the path name of a file exposes information about the structure of the local file system. To a lesser extent, the full path name of a registry key can provide the same information. We would argue that other resources are much less likely to betray sensitive information. Of the various Windows resource classes, only file names and registry key names contain potentially sensitive information.

We argue that first, little or no information need be collected for constant resources, since the resource can be located using invariant properties that can already be present in the profile.

---

[6] This underestimates the number of resources actually accessed. The Notepad execution included one or more additional threads whose execution began in a DLL. At least one thread managed the GUI; it required many more DLLs and other privileged operations. In the interest of focusing on the most important aspects of our problem, we omitted those threads from our analysis.

[7] It can be argued that `c:\Documents and Settings\All Users` should be considered a constant resource, as it is the value of the environment variable `%ALLUSERSPROFILE%`.

Second, little information is needed for variable resources, since particulars about a resource (such as pathname) are not relevant to the program and can be ignored. Third, it will often be possible to express invariant properties in terms of environment variables or the values of registry keys.

However, in some cases it may be desirable to collection information about some resources, for example to establish a relation between resources (such as several files located in the "home" directory of an application). We argue that this can be done without jeopardizing collaborator privacy, since our only concern during correlation is to check for the appearance of the same strings in path names. The key idea is to hash each element in a path name. The hashes can easily be compared for equality, while the original path names cannot be re-computed from the hashes.

Consider the following approach.

In the profile, each constant element is represented by a hash value. The profile might contain partial path names for constant file and registry key resources, as described above (see Section 2.1.2). This information is distributed to each collaborator site. The partial path names identify the constant elements of the pathnames and elide the variable elements. At collaborator sites, local file path names are hashed element-wise and compared with the (hashed) names in the profile. If a match is found, the file corresponds to a constant resource in the profile; there is no need to send any path name to the EPP central site.

Information about variable resources is also sent to collaborator sites. This information does not include path names, since such names are not meaningful. Instead it includes, for each variable resource, the set of program points (which implies a set of access rights) at which the variable is accessed. If a program resource is accessed at the same set of program points, it is considered an instance of the variable resource. Again, there is no need to send information about the resource to the EPP central site.

If the file access pattern does not match any constant or variable resource in the profile, the path name can be sent to the EPP central site, hashed element-wise, together with information about the program points at which the file was accessed. The element-wise hashing protects the path name from possible attempts to extract information about the local system.

4.2 Malicious and incorrect inputs

There is a danger that malicious collaborators will provide incorrect information, either to vandalize the system or to trick it into attributing excessive privilege to a target program.

Any defense against malicious collaborators must rely on the assumption that there are very few of them compared with the total number of collaborators. By requiring corroboration for new privilege requirements from alternative sites and users, we can avoid premature acceptance of spurious inputs.

What kind of incorrect inputs can occur? The malicious user can attempt to introduce new program points and new constant resources. He can also introduce new access rights on constant or variable resources. We treat each of these threats in turn.

First, the user can attempt to introduce a new program point that is not in the program. Introduction of new program counter values can be checked against the program binary. Introduction of new ntdll operations for a given program counter value is discussed below.

Second, the malicious user can introduce a new constant resource (for example, access to a protected system file or a confidential personnel file belonging to a user). If we exploit the fact that constants must appear in the program binary, then this threat is extremely limited.

Third, the user can attempt to introduce new access rights on a constant or variable resource. Because we identify resources by the set of program points that access them and required access rights are determined by the ntdll operation (part of the program point), this requires introduction of a new program point for the resource. To defend against this threat, we can require that each program point in the profile be corroborated by N collaborators, where N can be a small absolute number or a percentage of the number of collaborators. The former is probably preferable, since unusual exceptions may result in rare program points being reached.

In the limit, it may prove difficult to distinguish in the privilege profile between extremely rare—but correct and benign—uses of privilege and the effect of concerted attacks on the profile. This challenge should be addressed once privilege profiling technology is in place.

5 Discussion

We have shown that it is possible to collect information about a program's exercise of privilege from a variety of users and to produce an abstract profile of the program's use of resources. Such a profile could be used for program testing (avoiding unintended use of resources), program installation (establishing correct privileges for privileged programs), and intrusion detection (detecting anomalous use of resources).

In performing this work, we encountered instances of surprising exercise of privilege by Windows programs. For example, the commonplace Calculator writes to the protected file `C:\WINDOWS\win.ini` (to record use of Scientific mode) *only if* the user is an Administrator (users not in the Administrator group are not allowed to write to that file, so the write is aborted). The standard Windows file chooser, used by Notepad as well as many other applications, opens files with wild abandon for trivial reasons—for example, to obtain a custom icon to display next to the file name. While such discoveries are depressing, the use of tools such as we are working on would enable system designers and programmers to find ways to avoid them.

The basic premise of this work is that by observing a program's actual use of resources, we can find the least privilege that it needs. The profligate use of resources that the program does not in some sense *really* need threatens that premise. We hope to be able in the future to learn ways to distinguish vital from trivial use of resources—for example, by observing program behavior in restricted environments. At the very least, we can identify and expose the problems and contribute to their solution.

6 Acknowledgements

We owe the idea of creating an empirical privilege profile to Lee Badger, who suggested tracking a program's exercise of privilege at a large number of user sites, possibly in the manner of SETI@home [4]. The work described in this paper was supported under DARPA contract W3194Q-04-C-R260.

7 References

Information about privileges and ntdll operations is available from the Microsoft on-line library [5]. Microsoft does not publish detailed information about ntdll operations. We relied on a book that we and others have found to be reliable [6]. We also used [7], which is not as detailed as [6], but sometimes has additional information. We are also aware of other hacker resources [8, 9].

References cited in this report:

[1] Saltzer, J.H. and M.D. Schroeder, "The protection of information in computer systems," *Proceedings of the IEEE*, 1975. **9**(63): p. 1278-1308.

[2] Gao, D., M.K. Reiter, and D. Song. "On Gray-Box Program Tracking for Anomaly Detection," in *Proceedings of the USENIX Security Symposium*. 2004, pp. 103-118.

[3] Gao, D., M.K. Reiter, and D.X. Song. "Gray-box extraction of execution graphs for anomaly detection," in *Proceedings of the ACM Conference on Computer and Communications Security*. 2004, pp. 318-329.

[4] SETI@home, http://setiathome.ssl.berkeley.edu/.

[5] Microsoft Corporation, The MSDN Library, http://msdn.microsoft.com/library/default.asp?url=/library/en-us/secauthz/security/privileges.asp.

[6] Nebbett, G., *Windows NT/2000 Native API Reference*. 2000, Thousand Oaks, CA: New Riders Publishing.

[7] NTinternals.net, Undocumented Functions for Microsoft Windows NT/2000, http://undocumented.ntinternals.net/.

[8] Holy Father, Hooking Windows API - Technics [sic] of hooking API functions on Windows, www.assembly-journal.com/include/getdoc.php?id=244&article=157&mode=pdf.

[9] Ivanov, I., API hooking revealed, http://www.codeproject.com/system/hooksys.asp.