

Lecture Notes, 10/21/98

Secure Environment Java Security

I. Secure Environment

Approach: intercept all kernel crossings to create a secure environment

The problem:

- o how to run untrusted code
- o example: helper applications for netscape
- o goal: safety not correctness (the program can be broken, but it can't hurt anything as a result)

Why mobile code?

- o full-control over local UI
- o reduced traffic due to local decisions
- o between network protocols (complete end-to-end control)

Orthogonal security:

- o Key idea: guard only the entrances and exits
- o security must be orthogonal to functionality
 - security can't have bugs (bugs are security holes)
 - application code *will* have bugs
 - therefore, we must keep them separate (orthogonal)
- o traditional view: security must be integrated
 - this conflicts with orthogonality
 - better version: design should be integrated, code should not

Approach:

- o use Solaris debug facility to intercept all kernel calls (that matter)
- o check arguments, remove some calls
- o recursive contain child processes
- o use config file to control variable access such as file system and network

Base functionality:

- o process tracing facility allows interception of all kernel calls, turn on/off per call
- o not perfect, but good enough: e.g. can't always abort a system call, which means we must abort the whole application
- o can modify args and return values

Alternatives:

- o make each app secure: unrealistic and untrusted authors
- o firewall: can keep all or nothing out
- o java (discussed below): use JVM to provide security -- keep language secure rather than just kernel calls.

Policy modules:

- o control access explicitly
- o access denied by default
- o try to reuse policy modules to avoid configuration for each helper application

Example for helper apps:

- o forked children recursively traced
- o signals only self and children
- o limit environment variables
- o app placed in a subtree -- full access within, but can't get out
- o checks only on open, since read & write can only use open'd file descriptors
- o network access only to an X proxy (nested X system in a window; this prevents communication via X to other apps -- it's like you are the only X app running)
- o see samples from the paper

Apps:

- o ghostview, mpeg_play, xdvi, xv, xanim
- o recent additions: java and sendmail (!)
- o no real performance impact

II. Extensible Security Architectures for Java

IPC is too expensive to use for protecting separate subsystems... We need a way to make cross domain calls fast

- o Software Fault Isolation: add code to check to do bounds checking on every read and write. "play in your own sandbox only". Wild writes can still corrupt the program, but only its own code/data.
- o Type-based checking (java): ensure that "bad" pointers can't be created. (Implies: must GC (!), must do array bounds checking, must prevent casts (except to superclass,

or to subclass with dynamic check), no pointer arithmetic.) Assuming all of this, then multiple subsystems can exist in one address space, with no interference other than that enabled by the explicit creation of cross-domain pointers.

New challenges: now that we have more flexibility, how do we implement a reasonable security policy that is safe, extensible and flexible?

Three approaches:

- o capabilities: every pointer is a capability (hard to revoke)
- o stack introspection: enable/disable privileges a frame and its children, but disable when calling an untrusted class
- o name space management: change the class loader to use special versions of system libraries that limit and check access. Prevent any other way of using the real libraries by taking them out of the namespace. Kind of like capabilities, but completely not revocable.

Notes on language-based safety: can be done statically or dynamically

- o static => fast, but limited expressive power, since you can't do casting (even if safe).
- o dynamic => requires run-time checks, but need not prevent all things that could be unsafe (since you can check them explicitly at the last minute)
- o However, static prevents run-time errors and exceptions and thus has advantages for unsophisticated users and for program reliability.

Stack introspection:

- o privileges can't be global -- they only apply to one thread (so make them part of the thread state)
- o calling into untrusted code needs to disable privileges automatically (can't trust the callee to do it!).
- o need to automatically disable privileges when leaving the frame that enabled them -- this prevents a common but serious error: forgetting to disable privileges on ALL paths out of the procedure
- o Q: why is it bad to allow trusted code to access a resource if there is an untrusted ancestor on its call stack? A: because the untrusted code manipulate the trusted code in bad ways. For example, the "file open" code is trusted, but an untrusted procedure could ask it to open the password file.

Up-front target requests: it is a good idea to test all the privilege you might need at the start of the program!

- o Else, you get bad run-time errors when you try to complete an operation without permission (like saving your work!)