Foundations of Computer Security

Lecture 14: Covert Channels II

Dr. Bill Young
Department of Computer Sciences
University of Texas at Austin

Covert Channel

Definition: A *covert channel* is a path for the illegal flow of information between subjects within a system, utilizing system resources that were not designed to be used for inter-subject communication.

Note several features of this definition:

- Information flows in violation of the security metapolicy *though not necessarily in violation of the policy*.
- The flow is between subjects within the system; two human users talking over coffee is not a covert channel.
- The flow occurs via system resources (file attributes, flags, clocks, etc.) that were not intended as communication channels.

Covert Channel #1

Attempted access by S_L to a high level resource returns one of two error messages: `Resource not found` or `Access denied`. By modulating the status of the resource, S_H can send a bit of information on each access attempt by S_L.

This is called a covert *storage* channel because S_H is recording information within the system state.

Covert Channel #2

The KVM/370 operating system isolated processes on separate virtual machines. They shared the processor on a time-sliced basis. Processes alternated using the CPU, with each allowed t units of processing time. However, a process could relinquish the CPU early.

Process p could send a bit to process q by either using its total allocation or reliquishing the processor immediately. Process q reads the bit by consulting the system clock to see how much time has elapsed since it was last scheduled.

This is a covert *timing* channel because the information is recorded in the ordering or duration of events on the system.

Covert Channel #3

Processes p and q are not allowed to communicate, but they share access to a disk drive. The scanning algorithm services requests in the order of which cylinder is currently closest to the read head.

Process p either accesses cylinder 140 or 160. Process q requests accesses on cylinders 139 and 161. Thus, q receives values from 139 and then 161, or from 161 and then 139, depending on p's most recent read.

Is this a timing or storage channel? Neither? Both?

Covert Channel #4

An *implicit channel* is one that uses the control flow of a program. For example, consider the following program fragment:

```
h := h mod 2;
l := 0;
if h = 1 then l := 1 else skip;
```

The resulting value of l depends on the value of h.

There are sophisticated *language-based information flow tools* that check for these kinds of dependencies in programming languages.

Types of Covert Channels

It is possible to distinguish many types of covert channels, depending on the attribute manipulated:

- Timing: how much time did a computation take?
- Implicit: what control path does the program take?
- Termination: does a computation terminate?
- Probability: what is the distribution of system events?
- Resource exhaustion: is some resource depleted?
- Power: how much energy is consumed?

In practice, many researchers distinguish only *storage* and *timing* channels.

Lessons

- A covert channel is any path for information between subjects, utilizing system resources that were not designed to be used for inter-subject communication.
- A useful distinction is between storage and timing channels, though the breakdown is not always clear for specific channels.

Next lecture: Covert Channels III