Title: Audit logs: to keep or not to keep?

Topic: Practical considerations, specifically real-time v. post-mortem IDS

Speaker:

Christopher Wee
University of California, Davis
1 Shields Ave
Davis, Ca 95616
wee@cs.ucdavis.edu

Bio:
Christopher Wee is a Researcher in the Computer Security Group at the University of California, Davis, Computer Science Department. His current research interests are intrusion detection, security auditing and Formal policy specifications. Chris has participated in three intrusion detection systems projects: DIDS, a LAN intrusion detection system, GrIDS: a wide-area IDS, and LAFS: a file-system auditor.

In addition to his research, Christopher has developed control software for Cooperbiomedical, Technicon, Syva in biomedical divisions. And he has consulted for Northville Industries on automated commodities trading.

Subject category: real-time v. post-mortem IDS

Audit logs: to keep or not to keep?

Abstract

We approached this line of inquiry by questioning the conventional wisdom that audit logs are too large to be analyzed and must be reduced and filtered before the data can be analyzed or stored. The audit facilities of contemporary operating systems (Solaris, Windows NT) do not suit the needs of intrusion detection systems (IDS) well. Many types of computers (e.g., small, mobile, or embedded systems) do not have sufficient resources for conventional audit logging facilities. Our research proposes to create separate audit facilities to serve intrusion detection systems (IDS) and to meet the needs of long-term storage (archival) of audit logs.

In general, IDS want to characterize activity at the level of users, sessions and application transactions while audit logs present activity at the system call, process and network packet level of abstraction. Users do not want audit processing to detract from application performance. Thus, we are constructing audit processing modules in the kernel that perform pre-selection and reduction close to the source. This should significantly reduce the overhead of transferring large amounts of data across the kernel boundary and the wasted effort of generating audit data that is discarded by the IDS. The IDS specifies only the events and data that are reported, and none of the data are logged to disk. The principle is to monitor only that which we know how to analyze.

The large size of audit logs is not an impediment to long term storage and analysis given effective data compression algorithms. Solaris BSM and Windows NT audit logs are highly compressible using a Lempel-Ziv compression algorithm (i.e., gzip, pkzip), resulting in 80%-95% reduction in size. We investigate alternate audit log encodings that are compressed yet permit efficient retrieval, splicing and merges of incremental portions of the archived audit log. The audit log must also preserve supporting data that enable

us to understand the meaning of logged data (e.g., the UID to username mappings, or the security policy in force when the data are logged).