Logging Infrastructure and Log Analysis

Presented by Abe Singer

This tutorial will describe how to build an infrastructure to collect, preserve, and extract useful information from computer operating system and application logs -- ultimately to help the system administrator learn more about what is happening on their systems and network.

Logfiles hold a wealth of information, from resource utilization diagnostics to problems with hardware and software, security problems, and forensic traces of intrusions.

Many system administrators have been told to "go figure out those logs." It's a daunting task -- there's an awful lot of information in log files, unfortunately it's not well organized or codified. Formats of messages, even timestamps, vary between applications, and sometimes even between different versions of the same application; different operating system distributions will use different messages to record the same event; and the information you need may be spread out over several messages. The system administrator often ends up building a system based on the relatively random data seen early on, instead of having an idea of what they'd really like to know before starting.

This tutorial will show how to take a methodical approach to collecting and extracting information in an organized manner.

The focus will be primarily on UNIX syslog, with some discussion of Windows logging and other sources of log data.

Examples are heavily weighted toward security issues, but provide some examples of resource and diagnostic monitoring. Many real-world examples from logs are included throughout the presentation.

The tutorial consists of:

* What logs and log analysis are all about
* The overall process of building an infrastructure for log collection and analysis
* Identifying sources of log information
* Basic logging with syslog
* Centralized logging architectures
* Some alternatives to standard syslog
* Log management -- archiving, rotation, preservation
* Building an analysis infrastructure
* Simple analysis with basic Unix tools
* Preparing log data for analysis
* Log Reduction
* Log parsing, and parsing tools
* Log analysis techniques, statistical and other
* Windows logging, and forwarding to Unix syslog
* Legal considerations

Audience

People involved in system administration and system/network security, including forensic analysis. Some knowledge of Unix/Linux will be helpful.

Unix admins will learn about how to get information from logs for use in resource management, system monitoring, and troubleshooting.

Security admins will learn about using logs for intrusion detection, policy compliance, and forensic analysis.

People administering Windows-only networks will find the talk to be much more about Unix, but there will be discussion of windows logging, and analysis techniques covered are universal.

People who administer mixed Unix/Linux and Windows environments will also benefit from learning how to integrate Windows logs into a Unix logging environment.

Abe Singer

Mr Singer is a member of the Security Technologies Group at the San Diego Supercomputer Center. In providing operational security for the Center, he participates in incident response and forensics, and is expanding the SDSC logging infrastructure. His research is in pattern analysis of syslog data for data mining.

Mr. Singer is the author of “Building a Logging Infrastructure” (SAGE, 2004. http://www.sage.org), and is currently writing a book on Log Analysis, due out around January 2006.