

Detecting and Displaying Novel Computer Attacks with Macroscope¹

Robert K. Cunningham, Richard P. Lippmann, Seth E. Webster
MIT Lincoln Laboratory, Information Systems Technology Group
244 Wood Street, Lexington, MA 02420-9108

Purpose

Macroscope is a network-based intrusion detection system which detects and displays attacks that increase a computer user's privilege level. The technique extends Bottleneck Verification, which detects novel computer attacks by looking for users performing high privilege operations without passing through legal "bottleneck" checkpoints that grant those privileges. The extension allows Macroscope to detect intrusions that exploit trust relationships, and to detect trojan programs.

Method

A network sniffer carefully assembles packets into session transcripts that are then parsed using models of how legal command-line interactions between a user and a computer should behave. By parsing telnet and rlogin network sessions, and by detecting bottleneck violations and monitoring remote system access, attacks can be discovered without advance knowledge of the method of the attack.

Network traffic for each rlogin or telnet session is collected and the data from individual packets are assembled into two transcripts, one for each side of the duplex communication. While creating the transcripts, client and server control data for each protocol are removed (e.g., telnet option negotiation) so that only actual data are included.

The resulting transcripts are then parsed to detect and identify prompts, commands and their responses. Processes are marked to indicate if they are executing as a regular user or a super-user, and the transitions from one thread to another are examined for illegal transitions. For UNIX and Windows/NT systems, the only command that is authorized to effect the transition from a regular user to a super-user (i.e., root or administrator) is the `su` command. All other commands that effect such a change are an attack on the system. This style of context-dependent parsing reduces the amount of scanning required over previous keystring-matching detection systems that scan for a set of specific attacks, but the software required to perform this parsing is more complex. Macroscope knows about more than 270 unix and shell commands and has interpreters for 70 of these.

Results

Macroscope was used to process data from the 1998 DARPA Intrusion Detection evaluation, showing a false alarm rate more than two orders of magnitude lower than a reference key-string system, while simultaneously increasing the detection rate from roughly 20% to 80% for user-to-super-user attacks². For this corpus, Network Bottleneck Verification detected 79% of the attacks at the rate of about one false alarm per day. The system has also been used for off-line analysis of approximately a hundred and forty thousand internet sessions captured across three months at more than 100 different sites. In these data, 67 illegal transitions to root were detected with fewer than one false alarm per site per day. In several instances, the system detected multiple attacks in a single session, when an attacker broke into a single machine and then exploited network trust relationships to attack additional systems.

New aspects of work

Macroscope is a network-based intrusion detection system that detects attacks using an extension of the initial Bottleneck Verification algorithm. Whereas the initial implementation scanned for prompts and changes to prompts, the current implementation includes a model of a user's command-line interaction with a computer system, allowing it to detect attacks that span multiple systems. The increased sophistication improves the detection rate and substantially decreases the false alarm rate, allowing for accurate color-coding and indexing of session transcripts. These results are stored in a database and converted to HTML, for access via an SSL-enabled web browser.

Conclusions

Macroscope combines robust packet assembly, extensive session analysis, and clear display into a system that is exceptionally good at detecting attacks that advance the privilege of the user.

¹ This work was sponsored by the Department of the Air Force under Air Force contract F19628-95-C-0002. Opinions, interpretations, conclusions, and recommendations are those of the authors and are not necessarily endorsed by the United States Air Force.

² The DARPA results should be considered "unofficial," as some members of the intrusion detection team interacted with MIT Lincoln Laboratory's evaluation corpus development team.