

Using Mobile Agents for Analyzing Intrusion in Computer Networks*

Jay Aslam Marco Cremonini David Kotz Daniela Rus[†]

Department of Computer Science, Institute for Security Technology Studies
Dartmouth College
Hanover, NH 03755

Today hackers disguise their attacks by launching them from a set of compromised hosts distributed across the Internet. It is very difficult to defend against these attacks or to track down their origin.

Commercially available intrusion detection systems can signal the occurrence of limited known types of attacks. New types of attacks are launched regularly but these tools are not effective in detecting them. Human experts are still the key tool for identifying, tracking, and disabling new attacks. Often this involves experts from many organizations working together to share their observations, hypothesis, and attack signatures. Unfortunately, today these experts have few tools that help them to automate this process.

In this project we recognize that human experts will remain a critical part in the process of identifying, tracking and disabling computer attacks. We also recognize that an important part of the discovery, analysis, and defense against new distributed attacks is the cooperation that occurs between experts across different organizations. Many installations do not have the expertise necessary to develop full attack analyses. Our goal is to build automated tools for computer experts and system administrators to:

- identify the characteristics of an attack given data from network sensors
- develop a hypothesis about the nature and origin of the attack
- share that hypothesis with security managers from other sites
- test that hypothesis at those other sites and coordinate the results of testing
- archive the data necessary for use as evidence in later law-enforcement actions

*We are grateful to the DOJ for their generous support of this work.

[†]Corresponding author : rus@cs.dartmouth.edu

The main difficulties in catching the bad hackers arise because (1) existing system logs are too large; (2) existing system logs do not contain all the information that could be useful in tracking down an intruder; and (3) attacks often come from multiple sources and spawn processes at multiple destinations but there is no way of coordinating logs across administrative domains. We believe that mobile agent systems can play a key role in addressing these difficulties. Because formulating a correct hypothesis requires coordinating and correlating system logs from multiple locations, mobile agents are a well-suited processing paradigm. A mobile agent can travel to the location of the system logs, filter out the time-relevant or location-relevant information from these logs, and correlate all this information without necessitating costly file transfers across the net.

We envision a new generation of security systems that work across administrative domains and use mobile agent technology to share and correlate system logs. We believe it is feasible to organize groups of computers that have different administrative domains as logical entities and use a mobile agent system to access, share, and correlate its system logs. Such alliances are already becoming a reality. For example, the Information Technology Information Sharing and Analysis Center (IT-ISAC) is a consortium of 19 technology giants including Microsoft and Oracle who joined forces recently to share security data, to protect themselves from hackers by sharing reports of electronic threats, incidents, solutions and countermeasures.

The mobile agent security system can be built on top of any mobile agent system, such as our own D'Agents system. The main component of such a system would be (1) the basic agent infrastructure; (2) a module for the distributed capturing and accessing of security logs; (3) a module for correlating data from the security logs; and (4) a module for formulating at-

tack hypotheses based on the output of the correlation module.

We have already developed D'Agents, which we envision as the infrastructure for our proposed security system. We have also developed a module for the distributed capturing and accessing of security logs and more generally, of distributed information. We are in the process of developing the last two modules.

Serval, our scalable information-retrieval server, is based on mobile agents. This server is hosted on a cluster of Linux workstations, each running the D'Agents execution environment and each containing a portion of the document collection. Agents wishing to search the collection may jump to any one of the hosts, connect to the document index, and query the index for relevant documents. Then, they can choose to read the documents through an internal Network File System or by jumping to each host to examine the documents locally. The goal is to examine issues in constructing scalable mobile-agent services, balancing load among the hosts, and allowing the agents autonomy in deciding how best to service their query.

We plan to use Serval for intrusion detection by using the same infrastructure to store network and host log data. The database can index the logs for quick searching. Network administration tools can use mobile agents to examine the data collection efficiently and flexibly, because complex analyses can be encoded as mobile agents, sent to the Serval cluster, and examine the data with high-speed local access to the data.

We plan to further extend the Serval infrastructure beyond its specialized cluster. In a local-area network, intrusion-detection data is inherently collected at distributed sites, at the hosts or network segments where intrusions may occur. The log data is stored in local hosts (possibly specialized secure logging hosts), avoiding the need to transmit it to distant parts of the network. By installing the Serval software on these log hosts, we create a "virtual Serval cluster" as a distributed network service. Analysis tools can launch a mobile agent into any one of these log hosts, and the mobile agent can search the distributed index and jump to other log hosts as needed to examine the log data.

Consider a typical network intrusion. After a reconnaissance phase (i.e., the hacker scans a certain network to find hosts that exhibit a given vulnerable service), a vulnerability is exploited, for example executing a buffer overflow attack against a machine with ftp enabled. This allows the hacker to overcome the standard login process and gain root privileges on the compromised machine. The hacker, then, performs opera-

tions on the compromised machine (read/modify data, install backdoors) and tries to compromise further internal machines of the same organization. These further attacks are likely to be easier than the first one, since internal machines are often less protected than the network's perimeter. Trust relationships, shared file systems, weak passwords are examples of common internal weaknesses that an hacker could easily exploit.

Given this typical scenario, the Serval mobile agent infrastructure is useful for the following goals. In the incident handling phase, it can help identify all the operations performed by the attacker and their time sequence. Serval is necessary to completely restore the compromised systems and avoid further attacks. In a forensic analysis, it supports the correlation among logs from different machines and the retrieval of evidence of the attack distributed among different machines of the network, which is useful for proving the damages caused by the attack. If a response strategy is considered, the Serval infrastructure can be used to stop the intrusion at its early stages, by isolating compromised machines. Such a response is necessary since today it is unfeasible to assume that all attacks could be blocked at the network perimeter.

References

- [APR99] J. Aslam, K. Pelekhov, and D. Rus, "A practical clustering algorithm for static and dynamic information organization", *Proc. of the 1999 Symposium on Discrete Algorithms*.
- [KGNRCC] R. Gray, D. Kotz, S. Nog, D. Rus, S. Chawla, and G. Cybenko, "Agent Tcl: Targeting the needs of mobile computers" *IEEE Internet Computing*, 4(1) 58-68, July-August 1997.
- [RGK97c] D. Rus, R. Gray, and D. Kotz, "Transportable Information Agents", *Intelligent Information Systems*, vol 9. pp 215-238, 1997.
- [NN00] S. Northcutt, and J. Novak, *Network Intrusion Detection: An Analyst's Handbook*, 2nd Edition, New Riders Publishing. 2000.
- [NCFF01] S. Northcutt, M. Cooper, M. Fearnow, and K. Frederick, *Intrusion Signatures and Analysis*, New Riders Publishing, February 2001.
- [SMK00] J. Scambray, S. McClure, and G. Kurtz, *Hacking Exposed*, 2nd Edition, McGraw Hill, October 2000.
- [VK99] G. Vigna, and R. A. Kemmerer, "NetSTAT: A Network-based Intrusion Detection System", *J. of Computer Security*, 7(1), pp 37-71, 1999.