RESEARCH BRIEF

QUANTIFYING COMPUTER SECURITY

The challenge

Traditional approaches to computer network security validation have not been quantitative, focusing instead on specifying procedures that should be followed during the design of a system.

When quantitative methods have been used, they have typically been very formal, aiming to prove that certain security properties hold, given a specified set of assumptions. Or they have been informal, using a "red team" of experts skilled in the practice of security and with knowledge of the system being studied.

An alternative approach, which has received much less attention in the security community, is to try to quantify the behavior of an attacker and his impact on the ability of a system to provide certain security-related properties.

ISR's research assesses a specific attacker behavior—estimating the correlation between a scan and an attack—based on empirical data.

The potential

A scan is a reconnaissance technique in which an attacker checks for an exploitable target by trying to determine attributes of the target host, such as whether it is running, what services is it running, what the operating system is, and whether there are exploitable vulnerabilities.

The computer security community often states that scans should be considered as precursors to an attack. However, very few studies have been conducted to quantify the validity of this hypothesis.

If this hypothesis is correct, then port scans will be a good indicator for system administrators that an attack will follow. If the hypothesis is not correct, then port scans should not be seen as part of an attack.

The research

The researchers developed a testbed using target computers for monitoring attackers and collecting attack data. Various scripts were developed to filter and analyze the data, separating management traffic from malicious activity. Malicious traffic included port scans, vulnerability scans and attacks.

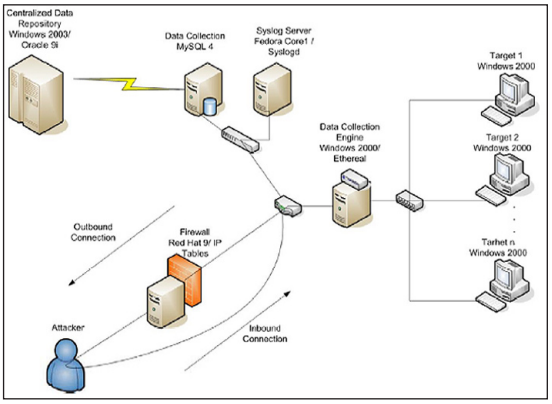

Testbed architecture

The filters separated the data into various scans and attacks directed to the target computers, based on the number of packets per connection. Access control, image control, event alerting, event logging data collection, data filtering and data correlation modules also were developed.

The correlation between scans and attacks was studied by first focusing on the scans and identifying whether attacks followed them, then analyzing the attacks and identifying the ones that had been preceded by a scan.

Three types of scans were considered. An ICMP scan uses the information provided by ICMP control messages to check the availability of a target machine and fingerprint the target operating system. This scan provides less information than a port or vulnerability scan.

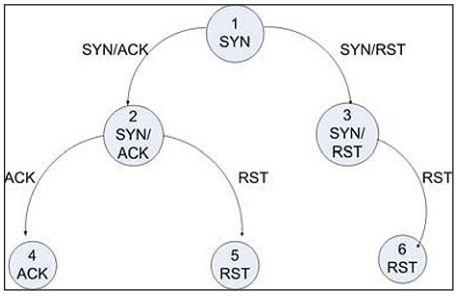

Representation of port scans

A port scan is used to check for open or closed ports and for used or unused services, which could have a vulnerability for the attacker to exploit. Moreover, the implementation of the TCP/IP stack is operating system-dependent,

so the attacker can use this information to fingerprint the target operating system.

Vulnerability scans check for specific vulnerabilities within specific services or applications. This kind of scan can be used to fingerprint the presence or absence of an exploitable vulnerability. The techniques to fingerprint different kinds of vulnerabilities vary, making it difficult to develop a generic algorithm to detect them.

In the research, almost no (.04%) ICMP scans on their own were followed by attacks. Contrary to the popular wisdom, port scans alone were followed by attacks only 4% of the time. Vulnerability scans alone were followed by an attack 21% of the time.

When all three scans were used, an attack followed 46 percent of the time; however only a small percentage of scans used all three methods. The best indicator that an attack would follow was the combination of a port and a vulnerability scan—attacks followed 71 percent of the time.

Significantly, more than half of attacks were not preceded by a scan of any kind. However, more than 38% of attacks were preceded by a vulnerability scan, either alone or in combination with other scans.

Results

Identifying port scans in combination with vulnerability scans launched from a specific source IP address is a good indicator that an attack will follow from the same IP address. However, port scans alone did not appear to be a good indicator of a future attack. Based on Dr. Cukier's research, port scans should not be considered as precursors to an attack.

Research team

Dr. Michel Cukier
Department of Mechanical Engineering
Center for Risk and Reliability
Institute for Systems Research affiliate
University of Maryland

Students
Susmit Panjwani, Stephanie Tan, Keith Jarrin
University of Maryland

Support

The University of Maryland's Institute for Systems Research IT unit and Office for Information Technology hosted and helped maintain the testbed for collecting the attack data. This research was supported by NSF CAREER Award 0237493.

Contact

Michel Cukier
Assistant Professor
Department of Mechanical Engineering
Center for Risk and Reliability
Institute for Systems Research affiliate
2100 Marie Mount Hall
University of Maryland
College Park, MD 20742

Phone: 301.314.2804
Fax: 301.314.9601
Email: mcukier@eng.umd.edu
www.enre.umd.edu/faculty/cukier.htm

Web links

Center for Risk and Reliability
http://www.enre.umd.edu/centers.htm

Reliability Engineering at the University of Maryland
www.enre.umd.edu/