Measurements and Security

Steven M. Bellovin
smb@research.att.com
http://www.research.att.com/~smb

At £rst blush, "measurements" and "security" do not seem to be commensurate topics. We have no idea how to measure the two most important aspects of security—the strength of a defender or an attacker—and it's quite possible that we never will. Nevertheless, there are important measurements that can and should be made.

The most obvious is the scope of an actual attack of some known type. Thus, "packet telescopes" [MVS01] can detect backscatter from certain kinds of distributed denial of service (DDoS) attacks. Other large-scale attacks have been studied as well, such as Code Red [MSkc02].

Though there are no current measures of the strength of an arbitrary attack, it is sometimes possible to estimate the strength of a particular threat. Consider cryptanalysis. A 1999 brute-force attempt to crack a DES key succeeded by harnessing the power of about 100,000 machines.[1] An enemy who controls a large "botnet" could do the same. How large are typical botnets? There are reliable claims that nets of up to 200,000 machines may exist, but there are no actual measurements of the size distribution.

Another class of measurement detects the incidence of vulnerable software. Rescorla, for example, measured the adoption of patched versions of SSL [Res03] after a vulnerability was discovered. Earlier, Cheswick conducted a similar survey of DNS patching [CBR03]. (They and others have found depressingly low repair rates.)

Measurements can also be used to estimate the cost of security measures such as encryption. During the debates over the design of IPsec [KA98], packet size distributions were used to calculate the approximate overhead.

Some security-related uses of measurement technology are less obvious. In [BG03], Bellovin and Gansner use topology maps as part of link-cutting attacks. But this attack, and more common forms of routing attacks, can be detected by measuring the change in paths and path lengths [RWXZ02].

Partridge et al. [PCJ+02] showed how timing measurements could be used for traf£c analysis on encrypted traf£c.

There are more things that we'd like to measure. Many forms of intrusion detection systems are, in fact, measurement and correlation devices. For example, one good indication of a new worm attack is a sudden increase in connection attempts to many different destination addresses.

In summary, much security—and attack—technology relies on measurements.

[1] http://www.distributed.net/des.

References

[BG03] Steven M. Bellovin and Emden R. Gansner. Using link cuts to attack Internet routing, 2003. Draft.

[CBR03] William R. Cheswick, Steven M. Bellovin, and Aviel D. Rubin. *Firewalls and Internet Security; Repelling the Wily Hacker*. Addison-Wesley, Reading, MA, 2003.

[KA98] S. A. Kent and R. Atkinson. Security architecture for the Internet protocol. RFC 2401, Internet Engineering Task Force, November 1998.

[MSkc02] David Moore, Colleen Shannon, and k claffy. Code-Red: A case study on the spread and victims of an Internet worm. In *Proceedings of the Internet Measurement Workshop*, 2002.

[MVS01] D. Moore, G. M. Voelker, and S. Savage. Inferring internet Denial-of-Service activity. In *Proceedings of the 10th USENIX Security Symposium*, pages 9–22, Washington, D.C., USA, 2001.

[PCJ+02] Craig Partridge, David Cousins, Alden W. Jackson, Rajesh Krishnan, Tushar Saxena, and W. Timothy Strayer. Using signal processing to analyze wireless data traf£c. In *Proceedings of the ACM workshop on Wireless security*, pages 67–76. ACM Press, 2002.

[Res03] Eric Rescorla. Security holes. . . who cares? In *Proceedings of the 12th USENIX Security Conference*, August 2003.

[RWXZ02] Jennifer Rexford, Jia Wang, Zhen Xiao, and Yin Zhang. BGP routing stability of popular destinations. In *Proceedings of the Internet Measurement Workshop*, 2002.