

On the Brittleness of Software and the Infeasibility of Security Metrics

How secure is a computer system? Bridges have a load limit, but it isn't determined (as "Calvin and Hobbes" would have it) by building an identical bridge and running trucks over it until it collapses.

In a more relevant vein, safes are rated for how long they'll resist

attack under given circumstances. Can we do the same for software?

I've sometimes quoted Lord Kelvin:

"If you can not measure it, you can not improve it."

"When you can measure what you are speaking about, and express it in numbers, you know something about it; but when you cannot measure it, when you cannot express it in numbers, your knowledge is of a meagre and unsatisfactory kind; it may be the beginning of knowledge, but you have scarcely in your thoughts advanced to the state of *science*, whatever the matter may be."

But I've reluctantly concluded that current architectures aren't amenable to metrics of the sort I want. Here's why.

It's well known that any single defense can fail. More precisely, we all know that any piece of software can be buggy, including security software—the list is alarmingly long. This means that whatever the defense, a single well-placed blow can shatter it. We can layer defenses, but once a layer is broken, the next layer

is exposed; it, of course, has the same problem. Paul Karger and Roger Schell noted more than 30 years ago the difficulty of defending against an attacker who had purchased a copy of a system and could probe each layer offline (<http://csrc.nist.gov/publications/history/karg74.pdf>). In most situations, we must defend against threats of precisely this nature. This is why metrics are so hard: the attacker's effort is linear (rather than exponential) in the number of layers, and each effort is low in cost.

What we need are defense systems with this exponential property—systems in which getting through two layers is proportional to the product (not the sum) of each layer's difficulty. If our defense systems have this property, we have some hope of measuring their strength. The constant for any one layer might remain small—this is, after all, software that we're dealing with—but we've somehow found a compositional principle that negates the linear effect.

Unfortunately, we don't know how to do this. One possibility might be to use randomization techniques to increase the attack constant for a particular layer, but we're still dealing with software, and the attacker might go around our randomization. Thus, we can use ran-

dom stack frame layouts, similar to OpenBSD's, to defend against buffer overflows; the attacker, though, could launch an SQL-injection attack. Perhaps we could use intrusion detection and repair at each layer. If we can do that, the holes won't stay open at first; the attacker will have to continually relaunch the attack. This presupposes that we can build such self-repairing software—research on it is still at a very early stage—and it won't be subject to the same brittleness. In any event, the repair would have to succeed before the next layer was penetrated or some autonomous attack code could continue the attack on inner layers, even though the outer layers had been repaired.

The problem, of course, is that such systems don't exist. Each layer's strength approximates zero, so adding them together doesn't help. We need layers of *assured* strength, but we don't have them. I thus very reluctantly conclude that security metrics are chimeras for the foreseeable future. We can develop probabilities of vulnerability, based on things like Microsoft's Relative Attack Surface Quotient, the effort expended in code audits, and the like, but we can't measure strength until we overcome brittleness. And until we can measure security, we can't improve it. □

Steven M. Bellovin is a professor of computer science at Columbia University. He received a BA from Columbia University and an MS and PhD in computer science from the University of North Carolina at Chapel Hill. He helped create netnews, or usenet news, and is coauthor of Firewalls and Internet Security: Repelling the Wily Hacker. Contact him via www.cs.columbia.edu/~smb.



STEVEN M. BELLOVIN
Columbia University