Security Is a Subset of Reliability

Reliability measures the deviation between the system and the specification. Security involves a subspace of reliability—only particular deviations—thus, security must be easier than reliability. Thumbs up. Hastening over the delicate premise that the specification is always accurate and up-to-date, we can roughly align security with the subset of reliability where the cost of deviation per unit time is very high. Thumbs down. This makes us wonder about measuring how risk tolerance scales and consequently where to point our thumbs. Some numbers follow....

Daniel E. Geer Jr.
In-Q-Tel

Daniel G. Conway
Augustana College

Security itself isn't cheap. Adi Shamir says that security and cost are inversely proportional: to halve your vulnerability, you have to double your expenditure.

Mandated reliable latencies in medical care are log scale:

Emergencies: within three minutes.
Urgent needs: within three shifts (24 hours).
Routine physicals: within three weeks.
Ratio of routine to emergency: 10,000-to-1.

Compare that to the digital world:

You can lose your machine to penetration in: 700 ms.
You can rebuild it in: two hours.
Ratio of rebuild to lose: 10,000-to-1.

The Meta Group, a well-known consulting group, says that downtime (unreliability) losses lie along a log scale:

INDUSTRY	Energy	Telecom	Manufacturing	Finance	IT	Insurance	Retail	Pharmaceuticals	Banking
REVENUE IN THOUSANDS OF US$ / HOUR	2,818	2,066	1,611	1,495	1,344	1,202	1,107	1,000	996

In fact, data reliability is already core to corporate survival:

Bankruptcy filings among companies that lose their data center for ≥10 days:
93% fail within the year.
Bankruptcy filings among companies that lose their data for ≥10 days:
50% fail immediately.

Yet real-world protections can yield perverse results. The US Federal Aviation Administration's one-decade impact estimate for requiring child safety seats on airplanes indicates that those seats would save five infants, but would cost 30 to 100 infants as families choose instead to travel by automobile.

And real-world protections can have an enormous dynamic range in cost:

Cost per year-of-life-saved for various public health measures:
Sickle cell screening for black newborns: $236
Heart transplants: $158,000
Seat belts on school buses: $2,800,000
Banning asbestos in automatic transmission components: $66,000,000

Luckily, security is generally not why systems fail:

Operations errors 60%
Applications 20%
Non-Security 15%
Security 5%
Ratio of error to treachery: 60%/5% = 12-to-1

But security does control whether redundancy is actually a help:

Where risks are uncorrelated, redundancy raises availability
Where risks are correlated, redundancy lowers availability

If we treated security-liability like safety-liability, then the liability threshold would be when $B < PL$, where P = the probability of loss, L = the amount of said loss, and B = the cost burden of adequate precautions.

The US Federal Trade Commission's 2003 identity theft data says that the 4.6% of the public that had an incident collectively spent 300 million hours and US$5 billion on clean up, or

$P = 4.6\%$
$L = 3 \times 10^8 hr \times \$5.15/hr + \$5 \times 10^9 / 10^7$ victims = $655/victim
$P * L = \$30.11 = B_{minimum}$

Ergo, is $30.11/consumer/year enough to prevent ID theft? If yes, then liability, else no liability.

Data protection reliability is tested by spotty data auditing.

INTERNAL	EXTERNAL	
	DONE	NOT
Done	24%	26%
Not	4%	46%

Ratio of those doing only internal data audit to those doing only external: 26%/4% = 4.25
Increase in total number of audits if everybody did both: 2.6 X.

Comparing real world reliability to digital security reliability:

Six nines (99.9999%): Mature manufacturing quality assurance
Five nines (99.999%): Public switched telephone network availability (it took 100 years to get there)
Four nines (99.99%): Domestic electrical transmission reliability
Three nines (99.9%): Maximum possible desktop uptime after downtime required by monthly patching drills
Two nines (99%): Credit-card number protection (640M issued vs. 112,205,775 reported exposed in 2007 and assuming 0.5% of exposed actually in play)
One nine (90%): Share of Internet backbone traffic that is not broadly related to attacks
Zero nines (10%): Aggregate ability of stock anti-virus to find new malware

We have a long way to go; Robert Frost may have said it best: "The woods are lovely, dark and deep/ But I have promises to keep/And miles to go before I sleep/And miles to go before I sleep." □

Daniel E. Geer Jr. *is the chief information security officer for In-Q-Tel. He was formerly vice president and chief scientist at Verdasys. Geer is a past president of the Usenix Association. Contact him at dan@geer.org.*

Daniel G. Conway *is an associate professor of business administration at Augustana College. He previously served on the business faculty at the University of Notre Dame and Indiana University. Conway's research interests include technology risk management and information economics. Contact him at danielconway@augustana.edu.*