Monopoly Considered Harmful

A recent position paper (www.ccianet.org/papers/cyberinsecurity.pdf) on software monocultures has generated heated discussions about whether monocultures exist and what they portend for us. Dan Geer and Dave Aucsmith address some of the issues

Daniel E. Geer Jr.
Consultant

The more advanced the society, the more it stands to lose from disruptions to its interdependencies. Only a few countries can experience electronic failure modes that could cause legitimate national security concerns. Those countries' fate is to make the mistakes that others learn from.

Two great risks

I was recently asked, on camera, what I thought the one or two biggest security risks were. If you're serious, this is a hard question to answer. Those who would color their answers with self-aggrandizement are cowards, if not knaves.

If this were math class, we'd start with axioms:

- More advanced societies are more interdependent.
- The ability to manipulate information is power.
- National security trumps more minor concerns.

Accordingly, only two kinds of risk truly matter. The first is the risk of effective attack on something that is fundamentally unique and essential—the approach radars at Reagan National Airport or the Fedwire or the GPS satellite array, for examples. Those who run such uniquely critical infrastructure are responsible for ensuring that the cost of attack is not worth it to the attackers. This generally involves defense in depth and replication; it is, basically, a referendum on the authority's willingness to spend money—critical, yes, but not our focus here.

The second critical risk—the one that concerns us here—is cascade failure. It is the forest fire of the networked world, the epidemic of the computer age. Cascade failure is, like calculus, the infinite sum of infinitesimals. Cascade failure fundamentally matters.

The risk of cascade failure is simple to envision: an attack on one computer turns it from victim into attacker, just as a person moves from infected to infectious. For such cascades to become national security risks, they need easy propagation. Ease of propagation peaks when all platforms are exactly alike: no decision-making is required, only the discovery of as-yet-unattacked systems.

Every computer security event of public note has been one of these two types; everything else has been a private tragedy. Those with the willingness to spend money on defense in depth and replication for unique assets have a solution. The real problem, which is without doubt growing more serious, is that of identical platforms that are riddled with security holes—the same security holes. It is silly to speak in euphemisms: the problem is with Microsoft's near-monopoly and the security characteristics of what that monopoly has wrought.

This is not Microsoft bashing; it is far more lethal because it is dispassionate. The identicality and flaw density in the Microsoft Windows monoculture present clear dangers to national security in proportion to the degree the nation in question depends on computers for the quiet enjoyment of its way of life.

Lessons in diversity

Nature teaches (those who will listen) that the richest ecosystems are the most diverse. Monocultures, to the extent that we humans insist on practicing them, require ever-increasing inputs of energy, fertilizer,

continued on p. 17

Dave Aucsmith responds: Diversity has a cost

Diversity has a cost. Enterprises have standardized on specific computer hardware and software to reduce procurement, operation, and maintenance costs. For many years, these costs were calculated as part of a computer system's "total cost of ownership (TCO)." These costs are real and are, like security, only one component of the cost-benefit analysis that every business or individual must compute.

Many enterprises have developed custom applications. Each computing base must develop specific applications. This would dramatically increase cost and greatly increase the environment's complexity.

To stretch the biological analogy in the CCIA report a little, organic systems contain biological diversity, which helps prevent failure induced by outside factors, such as pathogens. Yet, there will be a great commonality in the underlying biochemical mechanism because this mechanism has evolved as the most efficient method of performing a specific task (such as the Krebs cycle). Computer systems are similar. There is diversity in their defensive mechanisms, such as antivirus software, firewalls, and so on, and standardization in their interfaces to achieve operational efficiency. An airline does not try to maximize the types of aircraft it flies to reduce the chance of a design defect grounding the entire fleet. Rather, it standardizes to reduce its procurement, operation, and maintenance costs.

Security specialists easily forget that the actual goal of an enterprise is operating and maintaining computer systems so as to actually accomplish work.

Geer, continued from p. 14

pesticides, and surveillance. These inputs inevitably trend toward the diseconomic; one must, like the Red Queen, run faster and faster to stay in the same place. That moment of diseconomy can, of course, be delayed only so long as some otherwise-free good, such as clean water, can be absorbed as a hidden subsidy. In the case of computers, the hidden subsidy is the labor of systems administration.

The drumbeat is getting louder. Real data shows that the interval is decreasing between flaw discovery and flaw exploitation. The propagation rate between susceptible hosts regularly sets new speed records, and the total societal loss figures are growing steadily. Real data shows that Microsoft's near-monopoly on the desktop draws a near-monopoly of the attacks, in both type and number of attacks. Real data shows that once a computer is a victim, it is virtually certain to become an attacker. The percentage of total Internet traffic devoted to attacks is growing, and the half-life of a vulnerable machine is approaching zero.

Finally, real data shows what is absolutely national security writ large: the constant cacophony of amateur attacks is more than sufficient smoke screen for the real professionals to hone their craft as they wait for their moment on the world stage. Remember, the real measure of a virus writer's success is the labor cost of revisiting infected machines later. True professionals lie in wait for the next Nimda, with its multivector propagation and its new back door. When the time is right for them, they will send their chasers to exploit the machines that were alike enough for the virus to invade. And, after its passage, these machines will be identical. The only answer is platform diversity; without it you divide by zero. □

__Daniel E. Geer Jr.__'s research interests include a range of issues in digital security. He has an ScD in biostatistics from Harvard University's School of Public Health and a BS in electrical engineering and computer science from the Massachusetts Institute of Technology. Contact him at dan@geer.org.