Revolution through Competition?

Carl E. Landwehr
Editor in Chief

Competition, in the form of survival of the fittest, has been around far longer than humans, but humans are unique in setting up artificial competitions and awarding prizes for success. The Olympic Games is perhaps the oldest competition that continues to this day (albeit with a hiatus of two millennia), though no doubt it wasn't the first.

Technological competitions are more recent but not young. In 1419, for example, Fillipo Brunelleschi's innovative design won the competition to build the dome for the cathedral of Florence. In 1714, the British Parliament initiated a competition for an improved means of calculating the longitude of ships at sea, with the prize ultimately awarded nearly 60 years later to a technologist, clockmaker John Harrison, rather than to an astronomer, as many (in the Royal Society, at least) had wished. In the 20th century, prizes were successfully used to encourage practical aviation (the Orteig prize, which Charles Lindbergh won in 1927) as well as scientific progress in many fields (the Nobel prizes).

Today's scientists and technologists can compete for a variety of prizes that depend on advances in computing, sometimes in combination with other fields. The Turing prize recognizes basic contributions, but the Archon X Prize for Genomics (sequence 100 human genomes in 10 days), the Darpa autonomous vehicle Grand Challenges, the recently announced Google Lunar X Prize, and the more narrowly focused Netflix Prize are all attracting substantial attention.

Research funding agencies find that competitions are particularly attractive because they can open up a field to contributions from a wider variety of participants than the agency might normally reach, and they hold the potential for stimulating "free" research—something that the agency doesn't have to pay for out of its budget. A US$1 million prize might trigger donated research investments that would cost many times that amount.

A very real cybersecurity competition is in daily progress, between attackers and defenders of real systems, with real money at stake and real penalties for losers. Artificial cybersecurity competitions that have attracted the broadest interest have mirrored the real world through "Capture the Flag" (CTF) exercises, popularized by DEFCON and now at many other venues. These typically involve someone setting up a target, perhaps a network or a file to be defended, and the other side then trying to break down the defenses according to an agreed-upon set of rules.

In the Hot Topics in Security (HOTSEC) workshop at Usenix Security 2007, David Lie and M. Satyanarayanan proposed a method for quantifying system security based on open competitions mediated by a trusted organization. A product or system that succeeded in meeting the provider's security claims in the face of a public attack over a specified time period, motivated by an award of a given monetary size, would be issued a certificate validating this fact. In selecting products, purchasers could consider their certificated attack resistance.

Secure Computing took this sort of approach to demonstrate the strength of its Sidewinder firewall starting in 1995 with a US$10,000 prize. During a three-month trial in 2001, the reward increased to $100,000. Although no one ever claimed the prize, some say that the intensity of attacks mounted against the Internet-connected target had the effect of denying service to some network users in the area of the company's home offices in Minnesota.

The most important question to ask in setting up a prize or competition is, what are we trying to learn or achieve with it? If the goal is to teach students or system administrators how to organize networks securely in the current technological world, and how to respond to attacks in progress, the CTF model might be a good one. And if we're trying to learn something that's relatively easy to specify precisely, such as a new encryption algorithm, a competition could be a very effective means to do so. The US National Institute for Standards and Technology

(NIST) took this approach with great success, conducting a competition from 1997 to 2000 that led to the adoption of the Rijndael algorithm as the Advanced Encryption Standard.

But suppose we want an attack-resistant cyberinfrastructure that doesn't depend on hordes of well-trained administrators and cautious users to assure its integrity. Suppose we want an infrastructure whose attack resistance can be explained and understood, that's easy to manage and use, and that can be extended with new technology without requiring a complete re-evaluation of the system to assure that each new component doesn't compromise it.

Such a goal is much more complex and progress toward it more difficult to evaluate than is the development of a better cryptographic algorithm. One particular difficulty is that just as no car or airplane is built only to be "safe," no computer system is built only to be "secure." First is the requirement to perform some set of functions and then the requirement to do so "securely"—typically without compromising the availability, integrity, or confidentiality of information involved.

Perhaps a better model for such a contest would be the Solar Decathlon competition, sponsored by the US Department of Energy and related industry groups, which seeks to help students learn how to design and build homes that are both energy efficient and comfortable to live in. Each model home accepted into the competition is scored in 10 different categories. Some, such as energy balance, are strictly objective, but others, such as "architecture" can't be effectively measured. The point of the competition is to get students to think about these problems in new ways and to produce novel, innovative, yet practical solutions.

Could we undertake such a competition, a Cybersecurity Decathlon, to lead us out of the current cybersecurity mess? Suppose we identified some computing domains in which both functional and security properties could be reasonably clearly defined. We would still need to identify criteria (some objective, but for properties like usability and manageability, they could be subjective as well) for judging how well a particular entry met these properties.

Further, if the solutions developed as a result of a competition are to lead to stronger deployed infrastructures, industry will need to adopt them. Artificial competitions can impose bounds on competitors (such as disallowing social engineering attacks or limiting competitors to certain platforms) to focus advances on particular technologies or approaches. But effective solutions must eventually succeed in the relatively unrestricted competition of the marketplace. Incentives might be needed to involve industry experts from the start, either as sponsors or participants. Indeed, just formulating the competition properly could be a task for experts. If we can formulate the competition properly, the payoff could be revolutionary. □