DAN GEER

monoculture on the back of the envelope

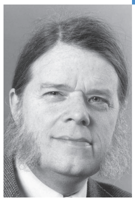

Dan Geer practices security medicine on corporate and government bodies of all sizes. For the privilege of doing so, he considers it a duty to report back whenever he is certain of what he has seen.

dan@geer.org

ABOUT TWO YEARS AGO, SEVEN various security folks released a paper where we tried to put together a single, coherent analysis [1] of the interaction of security and competition policy, which is to say, how a near-monopoly of Microsoft desktops affects the world's computing up to and including questions of national security. We weren't the first to use the word "monoculture" (that would probably be Prof. Stephanie Forrest at HotOS in Boston in 1997 [2]), and we invented nothing in the process; we were just the first to put it all in one place.

If you can call our paper a payload, the Computer and Communications Industry Association provided a launch vehicle, and at the last second my then employer showed up with a solid fuel booster already lit. We achieved orbit as measured by column inches in the global press and in many other ways as well—e.g., 10 days after our publication the CIO of the Department of Homeland Security was being grilled on the subject of monoculture on the floor of the House of Representatives [3], not that it dissuaded him from ending up with 200,000+ desktops, all Microsoft. Almost immediately, the NSF awarded Mike Reiter, CMU, and Stephanie Forrest, U. of New Mexico, a grant to study this very question . . . in the amount of $750,000 (http://www.scienceblog.com/community/older/archives/C/archsf373.html).

Finally, and as USENIX attendees will recall, there was even a formalized debate [4] on the question on June 30, 2004, at the Boston Annual Technical Conference.

Since then, has there been any great rush to diversify? No, even though the argument remains as valid as ever. There are exactly two paths to choose amongst with respect to monoculture security:

1. Embrace monoculture, since it allows you to get strongly consistent risk management exactly because everything is all alike.

or

2. Run from monoculture in the name of survivability.

Amongst the cognoscenti, you can see this: at security conferences of all sorts you'll find perhaps 30% of the assembled laptops are Mac OS X, and of the remaining Intel boxes, perhaps 50% (or 35% overall) are Linux variants. In other words, while security conferences are bad places to use a password in the clear over a wireless channel, there is approximately zero chance of cascade failure amongst the participants. Oddly enough, this exactly corresponds to Sean Gorman's work at George Mason, where he demonstrated a sharp turn for the worse when a single platform reaches 43% of the communicating total [5].

Statistics have been mounting up, of course, not that the existence of statistics automatically wins over hearts and minds (outside the cognoscenti, that is). For example, botnets assembled by automated means pretty much rely upon monocultured targets. Symantec's number is 30,000 added to botnets per day [6].

So, getting to the back of the envelope, what might just that number tell us? If 30,000/day is accurate, then we should be able to calculate the total infection percentage using total PC count, lifetime to repair/reload, and the 30,000 figure (which is technically "incidence" in public health terms) to get "prevalence" (the number currently diseased) and eventually to percentage. Doing that proverbial back of the envelope and blithely assuming static number of 200 times ten to the 6th PCs on the planet with 100 days between reloads or other forms of repair:

$$\frac{\frac{30 \times 10^4 \textit{ captured}}{\textit{day}} \times 100 \textit{ days} = 30 \times 10^6 \textit{ inventory}}{200 \times 10^6 \textit{ total PCs}}$$

Which gets you an estimate that perhaps 15% of all desktops are to some degree owned as I write this. This feels high, but as a personal data point, some colleagues recently found 70% of the desktops inside a defense contractor handling classified data to have spyware of one or another sort, and two keyloggers on the section head's desk. One can only assume that these are unusually careful folks, which thus reinforces the level of risk as high.

Let's look at cascade susceptibility terms but with an eye to the individual enterprise. As usual, there is an assumption, namely that when an infection enters the enterprise it will spread between and amongst those entities inside said enterprise. (This is what various people have called a "soft chewy interior.") Returning to the back of our envelope:

let: sizeof(enterprise) = y

and: Pr(individual_infection) = x

restated: Pr(no_individual_infection) = $1 - x$

hence: Pr(no_group_infection) = $(1 - x)^y$

Pr(group infection) = $1 - (1 - x)^y$

we want: LD50, that is x such that, given y, Pr(group_infection) = 50%

derivation: $.50 = 1 - (1 - x)^y$

$(1 - x)^y$ lineup $= .50$

$1 - x = .50^{1/y}$

$1 - .50^{1/y} = x$

(Notes: Pr = Probability of; LD50 = "Lethal Dose 50," the dose at which 50% of lab animals die.)

- For a 5,000 seat shop, there is a better than even chance of an attack taking down the enterprise when the risk of individual infection is .00014 (1 in 7,200) per user when integrated over the entire period of threat. For 100,000 seats, it's about 1 in 144,000 (.000007).
- That n(Web sites) ≈ 25,000,000 implies that each employee in that 5,000 seat enterprise must have an individual risk of infection less than 1 in 7,200; hence, for randomly selected Web sites, the density of infection must be less than 1 in 7,200: 25,000,000 / 7,200 ≈ 3,400, the number of Web sites that

can be infected across the entire Internet before a single random visit by each employee has a better than even chance of infecting the enterprise as a whole. For 100,000 seats, when n(infected Web sites) ≈ 175 for the entire Internet, a single random visit by each staff member has a greater than 50% chance of taking down the enterprise.

Summary

None of this is particularly good news but then again none of it is news at all. We knew this before, we just don't like hearing it, we shoot messengers, we try to patch things up. Everyone within the sound of my voice knows this. My 87-year-old cost accountant father knows this (his estimate is that over half of the productivity gains computers should have brought the domestic economy were lost due to standardization on the Redmond platform).

They know this in Redmond, too, where I do not envy the task they have in front of them, as it is like nothing so much as plugging shell holes below the waterline while under cannonade. In the meantime, Ballmer has one foot on the boat and one foot on the dock. The boat is labeled "Fix the security problem, but lose backward compatibility." The dock is the converse, "Preserve backward compatibility, but never fix the problem."

If he pulls his foot back onto the dock, he preserves backward compatibility but he never fixes the problem. This is betting that Microsoft is never tagged with liability for the security failures that only a monoculture can exhibit. Liability lawyers of the world are watching, and Steve is one nasty virus away from le deluge, not to mention the so-called progressive legislatures.

If he puts both feet in the boat and sails away from backward compatibility, then he absolutely puts into play the desktop in every single global corporation; those corporations are only sticking with Windows to amortize their existing investment in it. If they have to start over and write off that capitalization, they are not starting over with another round of "I won't hit you again, Honey, I promise."

And that, my friends, explains why Ballmer bought Connectix: the only way to introduce a new platform that arguably cures the security problem without kicking in the teeth of those who count on backward compatibility is to take the old insecure stuff and encapsulate it in some sort of virtual machine. It breaks the monoculture without breaking the monopoly, one part evil and one part brilliant.

REFERENCES

[1] D.E. Geer, C.P. Pfleeger, B. Schneier, J.S. Quarterman, P. Metzger, R. Bace, P. Gutmann, "Cyberinsecurity: The Cost of Monopoly—How the Dominance of Microsoft's Products Poses a Risk to Security," Computer and Communications Industry Association, September 24, 2003: http://www.ccianet.org/papers/cyberinsecurity.pdf.

[2] S. Forrest, A. Somayaji, and D. Ackley, "Building Diverse Computer Systems," *Proceedings of the 6th Workshop on Hot Topics in Operating Systems (HotOS VI),* May 5–6, 1997, p. 67.

[3] S. Waterman, "Homeland Security Software Eyed for Problems," United Press International: http://www.washingtontimes.com/business/20031020-092335-9325r.htm.

[4] D.E. Geer, S. Charney, A. Rubin, Debate: Is an Operating System Monoculture a Threat to Security?, Affirmative, USENIX Annual Technical Conference, Boston, Massachusetts, June 30, 2004.

[5] S. Gorman et al., "Is Microsoft a Threat to National Security? The Effect of Technology Monocultures on Critical Infrastructure," 2004: http://policy.gmu.edu/imp/research/Microsoft_Threat.pdf.

[6] J. Krim, "E-Mail Authentication Will Not End Spam, Panelists Say," *Washington Post*, November 11, 2004, p. E1: http://www.washingtonpost.com/wp-dyn/articles/A41460-2004Nov10.html.