

FUD: a plea for intolerance

Dinei Florêncio, Cormac Herley and Adam Shostack
Microsoft Corp., Redmond, WA

Even a casual observer of computer security must notice the prevalence of FUD: non-falsifiable claims that promote fear, uncertainty or doubt (FUD). We are bombarded with warnings of digital Pearl Harbors, the unstoppable of online hackers, and accounts of a cyber-crime problem that is said to rival the drug trade.

FUD sometimes masquerades as useful information though it is often “not even wrong,” in the sense of making no clear claim that can be checked: exact figures for undefined quantities, dollar estimates based on absurd methodology, and astonishing facts that are traceable to no accountable source. FUD provides a steady stream of factoids (e.g., raw number of malware samples, activity on underground markets, or the number of users who will hand over their password for a bar of chocolate) the effect of which is to persuade us that things are bad and constantly getting worse. While the exaggeration of threats hardly began with computer security, the field has certainly made FUD its own.

It may seem innocent enough to exaggerate in the service of getting people to take security more seriously; but we believe that reliance on factoids leads government and industry to spend wastefully and researchers to focus on the wrong questions. The scale of the FUD problem is enormous, and we argue that it prevents the establishment of security as a more scientific research discipline.

What's wrong with an illustrative story?

In offering information that is dubious, false or vague it creates avoidable confusion. Through creating the illusion that we understand when we do not and by injecting false facts, FUD oversimplifies complex questions, hindering our ability to grasp things that might actually be simple.

FUD makes it harder to form a coherent picture of the world. While the number of malware samples seems interesting, it says nothing about the success of that malware at infecting systems. Happenings on underground markets are certainly interesting, but activity doesn't translate into dollars at any fixed rate. Do we know if the passwords that people trade for chocolate are real or made-up? These details matter, and their absence hinders understanding. Much FUD comes in the form of factoids, which, of course, can be inconsistent, both with each other and with what else we know of the world. Who exactly lost a trillion dollars? Where are all the cybercrime billionaires? If cybercrime is so bad and people so careless why doesn't everyone have all of their money stolen every day? Purveyors of factoids make no effort to resolve these or other contradictions. Offered in isolation and selected for effect, FUD claims simply perpetuate uncertainty even on questions where clear answers might be possible.

FUD makes resource allocation difficult. Exaggeration might seem a harmless (or even necessary) tool: a short-cut to the right conclusion when the long way around seems too laborious. However, FUD doesn't just amplify, it distorts. When we inflate one threat we have to inflate others, as anything exaggerated 10x, seems small relative to things exaggerated 1000x. Nothing is so small that it can't be made to look enormous, nothing is so big that it can't be completely drowned out by the clamor to do something about some other inflated threat. We end up not being able to distinguish urban legends from real threats.

Finally, bad data drives out good. FUD supplies bad information in places where good might be possible. In addition to decreasing the signal to noise ratio, it makes it less likely that more scientific measurements will ever be carried out. If everyone already "knows" that baroque password policies and ninety-day expiration rules reduce harm, then it is less likely that experiments will test these claims. Once codified as a best practice and followed as a matter of course, the opportunity for observational experiments is limited. Unexamined assumptions thus remain unexamined, and errors that might have been caught persist. The hodge-podge of rules governing passwords have, until very recently, seen little serious research, but receive regular reinforcement in the form of factoids and anecdotes that confirm existing biases. Thus, FUD ensures, not merely that the true state of affairs is uncertain, but that it will remain so.

FUD is not a victimless crime

Those who produce good data suffer in a world that tolerates FUD. Good measurement work is possible [5,7] but expensive; those who produce it suffer if they must compete with FUD for attention. They are robbed of novelty if we feel “we already know” what the landscape looks like based on factoids and FUD. Doing high-quality studies is a losing proposition if those who consume information are indifferent to the care with which it is produced.

FUD erodes the ability to spend scarce resources sensibly and allocate effort where it will do most good. Indeed, in exaggerating uncertainty and danger, a main effect of FUD is to thwart the ability of defenders to spend efficiently. Thus everyone who wants good data to make decisions suffers because of FUD.

Finally, research suffers because of FUD. The professional credibility of computer security is impacted as ubiquitous FUD creates the impression of an unscientific field. Is cybercrime a trillion dollar problem and larger than the global drug trade? Is intellectual property theft ‘the greatest transfer of wealth in history’ or grossly overblown? Is cyber-war the ‘greatest existential threat we face’ or ‘the new yellowcake’?

It is difficult for a field to advance when unsubstantiated claims circulate unchallenged. The chances of placing security on a more scientific footing seem remote if we can’t agree what counts as knowledge. Are 99% of exploits really due to known vulnerabilities? Are 73% of compromises actually due to insiders? Explaining what has been observed and making predictions is always difficult, but it seems hopeless if we can’t agree on what is known. FUD obstructs the establishment of a more systematic approach to security problems.

Why are we telling you this?

Surely nobody argues for FUD. While security is awash in scare stories and exaggerations, members of the research community (and readers of CACM) certainly aren’t directly responsible for this state of affairs. However, as in many fields, “in order for evil to flourish it suffices that good men do nothing.” And flourish it certainly does: numbers whose quality is “below abysmal” get repeated by policy-makers [6] and trillion-dollar cybercrime numbers become the conventional wisdom. The answer to ‘How are we doing in security?’ is, to quote

Viega [4], ‘we have no clue.’ FUD could not achieve this without the acquiescence of many. It is this aspect that we wish to address.

Why is there so much FUD? Reuter [3] suggests that certain conditions favor the spread of fabulist claims and “mythical numbers”: the presence of a constituency interested in having the numbers high and the absence of a constituency interested either in having them low or accurate. Drug and crime statistics, for example, can easily become mythical: enforcement agencies have budgets that depend on the numbers being high, but no group is correspondingly interested in understating the numbers. The same one-sided bias applies in our domain. Unlike global warming, or the dangers of genetically-modified foods, where lobbies exist on both sides of the issue, many in security have the incentive to exaggerate dangers and few if any gain by understating them or ensuring accuracy.

The upward incentives may be beyond our ability to change. However, we can be the constituency that demands accuracy. In our roles as authors, practitioners and members of the research community we can put out the unwelcome mat for unsubstantiated claims. It is up to those of us in security research to avoid taking short-cuts in making the case for what we do. We can refuse to cite questionable reports, vague claims or outlandish dollar estimates. As PC members and reviewers we can ask our colleagues to aspire to a higher standard also. We don’t accept sloppy papers, so citing dubious claims (which are simply pointers to sloppy work) shouldn’t be acceptable either. An impressive collection of rationalizations is available to excuse the use of unreliable information: data is hard to get, cybercrime may be under-reported (how would one measure that?) and we do face active adversaries. However, unless we resist these temptations we look like a community that responds to uncertainty by lowering standards.

The broader computer science community also has a vital role to play. To have influence, FUD needs to spread. To spread widely it needs the efforts of many people. If they circulate unchallenged, after a while bad numbers come to be accepted as good; unless someone objects, the more they are repeated the more they become part of the consensus view. In this way numbers with little basis in reality [1,2] shape priorities and influence policy [5,6]. Anyone can help halt that process by refusing to forward, quote or tweet claims that don’t seem to add up. This needn’t be hard: FUD requires an un-skeptical audience and does not fare

well under scrutiny. Performing even basic sanity tests and asking ‘where did that number come from?’ or ‘how would you measure that?’ is often all it takes.

Security has many difficulties, but the field has no problem that FUD doesn’t make worse. It won’t just go away, but we can help stop the spread. We can make it more expensive to spread FUD than good information by challenging FUD-claims every time we hear them.

References

1. D. Florêncio, C. Herley, Sex, Lies and Cyber-Crime Surveys, Proc. WEIS 2011
2. A. Shostack and A. Stewart, New School Security, Addison-Wesley, 2008
3. P. Reuter, The (Continued) Vitality of Mythical Numbers, The Public Interest, 1987
4. J. Viega, Ten Years On, How Are We Doing? (Spoiler Alert: We Have No Clue), IEEE Security & Privacy magazine, Nov. 2012
5. R. Anderson et al, Measuring the Cost of Cybercrime, WEIS 2012
6. P. Maas and M. Rajagopalan, Does Cybercrime Really Cost \$1 Trillion? ProPublica, Aug. 1, 2012, <http://www.propublica.org/article/does-cybercrime-really-cost-1-trillion>
7. J. Bonneau, [The science of guessing: analyzing an anonymized corpus of 70 million passwords](#), IEEE Security & Privacy, 2012.