Breaking-in Research

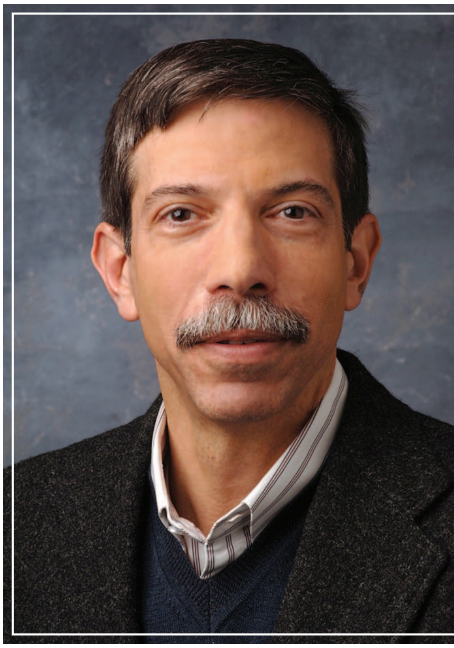

Fred B. Schneider
Associate Editor in Chief

Should the application of known attacks to compromise the security of actual systems be funded by research grants, as now seems to be happening? Should such work be eligible for inclusion in academic conferences or in the scientific literature? (There seems to be a trend in this direction.) Should a tried and true method of attack against yet another instance of some well-known vulnerability be considered *research* in the academic sense—work that is expected to remain relevant for a long time and concern a broad set of systems?

There are clear instances where retargeting extant types of attacks to today's systems does constitute research. These include

- work that shows the need for new kinds of defenses (and, ideally, new defenses would be proposed along with the attacks),
- work that illustrates new classes of vulnerabilities (perhaps due to new properties that need to be satisfied),
- work that enhances our understanding of the applicability for a class of attacks or defenses (acknowledging that researchers are not always diligent about describing applicability when proposing attacks or defenses), and
- work where the effort to retarget involves novel insight that itself will be valued as research.

So, there are certainly situations where attacking systems serves cybersecurity research in much the same way that building artifacts serves engineering research.

Even when it isn't research, *breaking-in*—identifying and reporting vulnerabilities—can be a valuable activity. Today's deployed systems are not secure. Moreover, our governments and other public institutions have failed to create incentives to change that reality. Stories in the popular press are one of the few ways for compelling software producers to fix vulnerabilities and for convincing them not to deploy new functionality that could be easily compromised.

Some technologies already have organizations whose mission is to evaluate artifacts for dangerous weaknesses. Underwriters Laboratories is a well-known example. But an organization like this doesn't exist for computer security. As a result, university professors join the ranks of consultants (who call themselves "researchers" but do not engage in what we call research) in the quest to uncover vulnerabilities that might be important to society but are not novel.

Although somebody does need to be uncovering these vulnerabilities, we should be careful not to portray the activity as "research." When we do, that mislabeling allows researchers and research funding to be diverted into activities that cannot move the field beyond its current reactive mode. Yet, absent an alternative engine for finding the vulnerabilities in today's systems, it would be socially irresponsible to ban such investigations by academic researchers. Research and vulnerability-finding are both important activities. So the real question is about incentives and funding.

A good starting point would be to clarify our expectations for academic researchers. First, we might revisit the mandate from research-funding agencies that a grant recipient not only undertake research but also engage in complementary activities to promote "broader impact" for the research. Why not accept the identification and reporting of vulnerabilities as a form of "broader impact"? This categorization would establish that breaking-in is subsidiary and complementary when it is not also research.

Second, public agencies whose mandate includes security for computing systems should take responsibility for funding faculty whose focus is on retargeting extant types of attacks to real systems. This would free research funds for use by those who really are engaged in research. We might even contemplate a further step, taking a cue from the professional schools. There, besides traditional faculty who engage in teaching and research, we find faculty whose primary responsibilities do not include the mandate to do research.

1540-7993/13/$31.00 © 2013 IEEE

Titles such as “clinical professor,” “professor of the practice,” and “extension faculty” are often given to those who fill such positions. If faculty dedicated to breaking-in (that is, not research) are moved out of positions where research is expected, then they could engage in their art with impunity. And others, who would do research, could be hired into the vacated traditional faculty slots.

Policies about what is accepted for publication and for presentation at our conferences provide another opportunity for creating the right incentives. Conferences do exist (Black Hat and DEFCON) for accounts of vulnerabilities and their exploitation. Academic research conferences could defer to such venues and simply reject submissions that describe the deployment of known attacks for compromising the security of actual systems. Or, if that’s too drastic to imagine, then our academic research conferences should label as such any papers that were accepted not for their research contribution but for the broader impact of exposing vulnerabilities in actual systems.

Great research, by definition, will have valuable impacts. But just because an activity is undertaken by a researcher and has valuable impacts does not make it great research—or even research. Breaking-in might have valuable impacts, but it isn’t always research. And the academic research community needs to adjust for that reality. ■

Acknowledgments

Although the opinions in this editorial are my own, I benefited enormously from the discussions during a breakout session I chaired on this subject at the November 2012 NSF Secure and Trustworthy Cyberspace Principal Investigators Meeting in Washington, DC.

Selected CS articles and columns are also available for free at http://ComputingNow.computer.org.