answer would give actionable data: "to reduce the odds of a successful insider attack, we need to reduce the number of employees with privileged access." The best answers, though, say what to do in the context of the actual operating environment. Why do so many people need to have privileged access? If you can't change that, what is the fallback position?

There are two obvious problems with this solution. First, giving usefully correct answers requires deep expertise, and that's in short supply. Second, even experts may not have the information they really need to give a deep answer. That in turn points to two challenges: we need education and training programs to develop deep expertise, expertise that goes beyond ever-longer checklists. And we need a deeper technical understanding of what makes systems—not individual programs or computers, but *systems*—secure. ■

Steven M. Bellovin is a professor of computer science and affiliate law faculty at Columbia University. Contact him via https:// www.cs.columbia.edu/~smb.

Digital Object Identifier 10.1109/MSEC.2022.3211758

Steven M. Bellovin
Columbia University

What Do We Owe?

If you're reading this, you probably work in computer ("cyber") security and/or privacy, whether as a researcher or practitioner. That in turn means that there are many stakeholders for your work. What do we owe them? What can we actually pay?

There's no shortage of products—firewalls, encryptors, intrusion detection systems code analyzers, monitoring services, and more—that purport to make our systems and sites safer. Vendors ship more secure base systems than ever before. That said, there are new reports of code vulnerabilities, ransomware use, supply chain attacks, etc., almost every day. But as the old saying goes, there's many a slip 'twixt the cup and the lip. What does this imply, as an ethical matter?

My goal here is not to analyze the many causes of security failures; many other people have done that. Rather, I want to talk about the ethical implications. Do we promise more than we can deliver? Why? And what should we do about it?

A technical piece of the puzzle is that many security failures happen not from deploying a single product, but rather from how they're deployed, in a complex network. That is, the security of a collection of single devices is not the same as the security of any one of the devices. Add to that the myriad configuration options, not all of which are compatible in a security sense with options on other devices. We do not know how to understand the security of such a setup; worse yet, we don't know how to talk about it. There are rules of thumb ("Keep systems patched. Limit privileged access.") but they'll only take us so far.

The security of a collection of single devices is not the same as the security of any one of the devices.

Consider the case of the ordinary password. Every security expert will rightly opine that passwords are a very bad way to authenticate to web sites; some form of multifactor authentication should be used. But we often rely on passwords to decrypt our disks at boot time. Is that wrong? What if the computer being booted is a high-security system? In other words, whether passwords are a good idea or a bad one is context-dependent.

Of course, even individual products can be flawed. Microsoft's products are, in general, excellent from a security perspective, but almost every Patch Tuesday includes fixes for critical vulnerabilities. Should there be a security warning? What should it say? "Abandon hope, all ye who hit Enter here" is probably a non-starter, but the fine print in most end-user license agreements says more or less that. What should we say?

One starting point is the ACM/IEEE Software Engineering Code of Ethics: "1.06. Be fair and avoid deception in all statements, particularly public ones, concerning software or related documents, methods and tools." That's good as far as it goes, but it doesn't go far enough. No possible statement can address all conceivable uses and configurations of products, especially if in combination with other products.

What, then, can we say? Clearly, we have a duty of honesty: honesty to our colleagues, to management, and ultimately to the public—but what is honesty here?

The trivial solution, to say "this is potentially insecure" in any situation, is correct but useless. People need more nuanced answers: "this can probably resist some outside hackers but not the better ones", or "an intelligence agency can likely break in easily." A better

Digital Object Identifier 10.1109/MSEC.2022.3205512
Date of current version: 28 October 2022

continued on p. 83