More is Not the Answer

Cormac Herley
Microsoft Research
One Microsoft Way
Redmond, WA, USA
cormac@microsoft.com

1. THE MESS WE'RE IN

Usable security has progressed a great deal in the last ten years. Initially, simply establishing the importance of the human factor in security was an uphill battle. That the interface of a crypto product, rather than it's users, might be to blame for it's lack of adoption [3] was not the accepted wisdom in 1999. That security might need to learn about users [1], as well as the other way around, was far from the popular view. Things have changed much. That the capabilities and understanding of users need to be factored into the design of security technologies is no longer controversial.

Users are the notoriously weakest link in many security chains; the easiest way to gain access to many resources is simply to convince, trick, or spoof someone into granting it. The attack that compromised RSA's master keys began with an employee clicking on a malicious link. Stuxnet bridged the airgap between isolated machines by correctly assuming that people would ferry files between them using USB thumb-drives.

A major disappointment then for many security professionals is the continuing lack of engagement on the part of the computer-using population. That consequences can be serious, and that much of the risk is related to user behavior seems not in doubt. Yet, the online population seems largely impervious to efforts to get them to take security more seriously.

An approach to explaining this fact is offered by Beautement *et al.* [2], who suggest that users have a limited tolerance, or budget, for compliance with security requests. Beyond a certain threshold increasing demands are simply met with attempts to circumvent onerous procedures. The thresholds appear to have been long exceeded for most users. In previous work [8] we suggested that from a cost-benefit standpoint users are rational to reject much security advice: the burden imposed is simply too great for the benefit received. In other words we might view the lack of care that users show for security as disappointing. A more extreme view might consider it dangerous or even a threat to national security. However, it is certainly no puzzle: if users reject the bargain they are offered it is simply because we have failed to make the case in a convincing fashion.

If we are unhappy with this state of affairs (and it appears that many are) then change is needed. The emphasis on studying user behavior is a welcome improvement. However, it is tempting to believe that human factors can be used as a toolbox of techniques to increase the time and effort users spend on security. For example, perhaps better design of security indicators and warnings might increase the notice people take of them, and better password strength meters might nudge them in the direction of greater strength. While these are interesting approaches, we argue that this still ignores the fundamental problem: spending more time on security is not an inherent good. Before asking for more of anything we should demonstrate that we're making good use of what we've got. Not only are we unable to demonstrate this, but it has become increasingly obvious that we've used our resources wastefully. We spend too much time asking for more and not enough optimizing use of what we've got.

In this paper we explore why progress has been slow and examine several possible directions. First, the scale and diversity of the web makes one-size fits all approaches hard. Second, the competition for user attention is fierce: there are no pools of unexploited user effort to be had. Third, persuasion is the only tool we have, mandates being often impossible or undesirable. Charting a way forward in these circumstances is hard. However, we outline several steps to improve the field.

1.1 Scale and diversity of the modern web

It is generally considered that there are upwards of two billion Internet users in 2013, while the total number of people with access to a networked computers in 1990 was about 2 million. Not only are more people using computers, but they are using computers more. The networked machines of 1990 were expensive devices, housed in controlled-access environments and were used generally for high-value tasks. Today, by contrast, hundreds of millions carry smart phones, laptops and tablets. Instead of a single computer account, many in the developed world will have on the order of one hundred

password-protected accounts, covering everything from banking to email, online video streaming to social networking. Facebook reported having 1.15 billion active users in March 2013, and numerous services have hundreds of millions.

Scale raises the cost of any changes we might propose. Consider the case of passwords. If we assume that each of two billion users spends 5 seconds a day typing passwords, it amounts to $2 \times 10^9 \times 5/(60 \times 60 \times 40 \times 50) = 1,389$ man years per day of human effort. Any technology, for example, the seeks to replace passwords with a technique requiring more than 5 seconds imposes an enormous aggregate burden. As we point out in previous work [8] an hour from each of 180 million online users (in the US) is worth approximately $2.5 billion. A major error in security thinking has been to treat as free a resource that is actually extremely valuable. Thus the importance if getting things right and not using user effort wastefully cannot be overstated.

As impressive as the scale of the web is its diversity. Not only is there enormous diversity of web services there is enormous diversity in how those services are used. The same service might be of great consequence to one user and almost unimportant to another; *e.g.*, an investment site may have accounts with portfolios of millions of dollars and ones with a few hundred, a social networking site may be all-important to one user and all-but unused by another. The same is true of hardware: some computers are used simply for surfing the web, while others are used to handle major financial transactions or control critical infrastructure. Thus, essentially the same security technologies end up having to protect assets that vary in value by orders of magnitude. The same browser and the same operating system (and the same sets of security warnings and cues) handle a great range of assets. The same commodity hardware and operating systems control nuclear centrifuges and are used for email and surfing the web. It is difficult to design mechanisms that are equally appropriate for assets that differ in value by orders of magnitude. Something that is up to the task of protecting high-value assets and critical infrastructure will be wasteful and unnecessary for low-value assets. The level of user care that can be expected varies by orders of magnitude. It is extremely difficult to design advice and usable security mechanisms that are divorced of context.

1.2 Competition for user attention is fierce

The answer to many usable security questions seems to be "more". We demand more password strength from users and ask them to change them more often. They should pay more attention to errors and warnings, and stop and think more before they click. They should educate themselves about the safety and security indicators of their browsers and operating systems and take more notice when anything strange happens. They should spend more time considering the potential consequences of their actions.

In pondering the hardness of making this happen it is worth remembering that competition for user attention is fierce. Most obviously, millions of web-sites, blogs and feeds compete for notice. A majority of the most popular sites are advertising-supported, where there is a direct relation between traffic and dollars. Search results are surrounded by sponsored links. News stories with "clickbait" headlines are covered in banners and preceded and followed by offers to take surveys. Articles are broken into several pages to multiply the number of advertising impressions served. Advertising popups drift across the screen at a speed calculated to make hitting the "close" button especially difficult.

The battle for user attention is not limited to advertising. Installing a new application is a process often peppered with interruptions about installing unrelated applications, receiving updates, placing icons on the desktop and changing the default search engine. New machines often come laden with free trials of software such as anti-virus packages, which noisily announce their presence and make increasingly insistent pleas as the end of the trial period approaches.

The reason for this theater of distractions is simple: user traffic, attention and eyeballs are the coin of the realm online. Most successful web companies make money by monetizing their access to users. This complicates the question of security for two reasons. First, security indicators appear, and security decisions must be made, with a carnival of distractions running in the background. Even if security indicators were clear, consistent and unambiguous (and they most certainly are not) they are hard to notice against the constant attempts to lure user attention elsewhere. Second, in asking for user time and attention we are asking for the single thing that is in most demand on the Internet. Facebook was valued at $100 billion based on the expectation that the company can translate their enormous share of user attention and interaction into earnings.

Thus, security must make its way in an extremely competitive environment. Not only are there no unclaimed pools of user effort to be had, it is difficult to preserve existing pools from incursions. It is hard to reserve time, effort, screen real-estate or techniques for security when each of them is a valuable and monetizable resource.

1.3 Persuasion rather than mandates

It is on ongoing frustration for many that users show so much reluctance on security matters. A natural question then, is why so many things are optional rather than mandated? If the correct course of action is clear and straightforward why is the incorrect one even an

option? If weak passwords are such a threat why do so many sites allow them? If the right action on receiving a warning is obvious why is any other course even offered?

The answer, of course, is that the correct course of action is seldom as clear as we might like. Many things can't be mandated for the simple reason what we are asking the user to do is not clearly defined; and when we substitute something that is clearly defined it doesn't accomplish the original goal. Onerous password policies can be mandated, but this doesn't guarantee good resistance to guessing attacks [10]. Password expiration can be enforced but, again, this does little to enhance strength [15]. While we can ask users to pay attention to security warnings we can't anticipate whether any particular warning will be a true or false positive. We resort to asking users to avoid "suspicious links," because we have nothing approaching a clearer definition. Thus, things are often not clear and we must rely on the judgement of the user. That being the case, we must persuade them that it is worth their while.

We've seen that the contention for user attention is fierce. The burden of proof for those who ask for more of it is thus large: that the benefit is greater than the cost must be shown, not assumed or asserted. In the consumer space it is users who bear the cost of increased security measures and users who receive the benefit. It is also users who make the decisions about where to make an effort and where to take shortcuts. Clearly, the perceived benefit to users is less than the perceived cost. While it is possible of course that they are wrong (*i.e.*, that the actual benefit to users is greater than the actual cost, even though they don't perceive it that way) this hardly matters: if users cannot be forced into measures then their judgement of the benefit and cost is the final court of appeal in the matter. Complaints that they don't understand, are mis-informed and so on, are simply distractions from the fact that the evidence as presented doesn't convince those who decide.

While we frequently resort to worst-case analysis and scenarios we seldom provide clear evidence that security measures reduce harm to a degree that merits the effort. Since competition for access to users is fierce and mandates work only in limited settings, like everything else security must make the case for the resources it wishes to consume.

2. PROTECTING USERS LESS BADLY

While usable security has had many successes in pointing out the failings of security UI, progress has been slower at providing actionable alternatives. It is difficult to give prescriptive answers on how things might be done better. Password advice is bad [11], but how might we do better? Security indicators are ignored [13], but what should we use in their place? Suspicious links may be hard to define [8], but we can hardly just ignore the problem? While the search for clear steps to do well is hard, a more modest goal is to do less badly. We next explore a number of approaches.

2.1 Never give an order you know will be disobeyed

As we've seen above, the answer to many security questions can seem to be to ask more of users. Even if we find their reluctance disappointing it can no longer be considered a surprise. (In fact, it is argued above and elsewhere [8, 11] that users are right to reject a bargain that offers a poor cost-benefit tradeoff.) That is, we have known for some time that users persistently seek short cuts, and complying with security requirements appears to be low on their list of priorities. This isn't surprising since security is seldom the main task, the benefit received is seldom salient (and is almost never shown to be greater than the cost), and there are many competing (and more compelling) demands on their time.

It seems safe to assume that this will continue: overwhelmed users will do the minimum on what is mandated, and ignore what is optional. Since we know this, plausible deniability is gone. We can no longer feign surprise that passwords are widely re-used and popup warnings ignored. It follows that observed user behavior must be considered a constraint, and realistic security designs mustn't assume more. Security regimes that assume higher levels seem destined for predictable failure.

Yet, the security advice offered to users by security experts, service providers and government agencies is filled with advice that we know has no possibility of being followed. Some of it is just unworkable; it imposes a burden that no reasonable user can pay. Some of it is just too vague, and does little beyond confuse the issue. A lot of advice has poor cost-benefit tradeoff and is overkill for the assets at stake. Finally, there is just too much of it and the cumulative effect is overwhelming. Without guidance on what to respect and what to ignore users are left to their own devices.

This all has an effect on credibility. If we insist on the necessity of measures that are ignored wholesale we simply draw attention to that gap between what we consider necessary and what users find they can get away with. It suggests that our goal in giving advice is something other than reducing harm.

2.2 When you don't know say you don't know

Confessing ignorance can seem incompatible with being considered an expert. Yet, security claims often appear little better (and in some cases much worse) than guesswork. Honesty demands that we be rather more frank about the limits of our knowledge.

On the question of password strength much mischief

has been caused by our unwillingness to admit that we don't have a clear understanding of how to measure strength, how to achieve it, or how much of it is needed. This has led to our insistence on the importance of measures that turn out to be almost unrelated to guessing resistance [10]. It leads to password meters which classify "Pa$$w0rd" as "very strong", and "wpnfusg" as "weak." It has led to the knock-on effect of countless organisations mandating complex password policies that appear unrelated to security [1, 7], and for periodic password changes that appear to accomplish little [15]. As we argued in Section 1.2, the scale of deployment is such that inefficiencies are a luxury we cannot afford. Our share of user time and effort is too valuable to waste; measures where there is no compelling evidence of efficacy should be reconsidered or dropped.

As outlined in Section 1.3 a great deal of what we ask of users can't be mandated (or implemented in the OS or browser) because it is vague. As pointed out in Section 1.1, the web is a diverse place. Any rule that we can think of as to how URLs should look will almost certainly have many benign exceptions. In longtail phenomena treating anomalies as "suspicious" is asking to be flooded with false positives. If we cannot describe clearly what types of links users should avoid it is better to be silent than offer frustrating and ambiguous instructions.

The frequency with which harm happens is important and a matter on which we are largely ignorant. Bad things certainly happen; people have passwords snooped, machines compromised and money stolen. Negligence on security matters can and does lead to real grief. However, it is not the case that a weak password always leads to theft, or that clicking through a warning always leads to compromise. A major point of disconnect between security practitioners and users appears to be the difference between worst-case and average outcomes [8]. If we lack firm evidence of average-case, honesty demands that we admit it rather than invoke worst-case harms.

Finally, estimates of cybercrime losses are often invoked in an effort to sell security measures. Reports that cybercrime is bigger than the global drug trade and that Identity Theft is rapidly growing can seem convenient props to help make the case. Here great caution is needed. First, it doesn't seem effective: if exaggerated claims were a useful tool in influencing users toward more security it would have worked by now. Second, estimates of cybercrime losses are notoriously bad and many turn out to be generated using unsound statistical methods [6]. Resorting to tainted evidence to sell the importance of what we do suggests an inability to convince by honest means.

2.3 Don't deny the obvious.

It is convenient in security to abstract all context away and consider only technical measures. This allows us to ignore questions of gain, cost, loss and motivation. While convenient, and a very natural technical approach, this fails to differentiate between high- and low-value targets. The passwords that protect a substantial bank account and a throwaway email account need not be treated with the same care. Sometimes, of course, low-assurance resources can be leveraged into high-assurance ones. However, to claim that this generally happens is to deny the obvious and ignore a major inconsistency: if worst-case outcomes are typical then why isn't everyone hacked every day? Most security advice errs on the side of caution, and is appropriate for high-value assets, where unbounded attacker effort must be expected. However most users have many low-value assets where attacker effort that is greater than the expected value is unlikely. Thus, we have a large void. We have some understanding (albeit imperfect) of the measures necessary to protect high-value assets, but we lack good tools to decided which of those measures can be neglected when protecting low-value ones. The burgeoning field of security economics [12] holds some promise in this direction, but much more work is needed here. However, pretending that low-assurance problems must be treated with the same care as high-assurance ones is counter-productive and puts us in the awkward position of insisting on the truth of something that billions can see for themselves is false.

We cautioned above against using worst-case scenarios when average-case is what users care about. A related error is to exaggerate the frequency of harm or average losses. Generally things are not as bad as we say. The Internet has two billion users; mostly they derive more good from it than harm. Bad things certainly happen. A large collection of would-be criminals seek to prey on the online population. However, turning code and stolen passwords into money is a lot harder than it looks [5]. It requires an almost wilful ability to ignore inconsistencies to espouse a view of the world in which ordinary consumers regularly lose money to cybercrime.

2.4 Be prepared to admit mistakes

The contention for user time, and the scale of the modern web would argue that user time and attention be used only as a last resort. We should ask for it only after exhausting all other possibilities. Unfortunately, the evidence suggests that we have used user effort as a first resort, not last.

For example, passwords have long suffered from the problem of off-line guessing attacks. This has given rise to a large variety of policies (which constrain the composition and length), expiration rules (that force them to be changed regularly), and tips and advice (that govern their choice and maintenance). These are all tools

to address the end problem, which is off-line attacks on passwords (although as noted earlier their efficacy seems much lower than generally assumed). Of course, many engineering problems have more than one solution. Trying to alter user behavior is certainly one way addressing off-line guessing attacks; however, it is by no means the only way. Elegant back-end solutions exist which make off-line attacks no worse than online ones [4]. Viewed in this way, addressing off-line guessing by asking effort of users is an $\mathcal{O}(N)$ solution to the problem (where N is the number of users); the solution proposed by Crescenzo *et al.* while not totally independent of N, might be $\mathcal{O}(\log N)$, at worst. In earlier times, when the number of users was orders of magnitude smaller, some inefficiency may have been tolerable; however, consuming $\mathcal{O}(N)$ resources on problems for which efficient solutions exist is no longer tolerable.

Indeed the whole question of addressing off-line attacks with user effort seems misguided. Encouraging users to choose passwords that will withstand online attacks seems relatively easy, and a good use of effort. Trying to get them get them to devote the additional effort of withstanding off-line attack is extremely hard, and largely futile. The difference between the strength of password needed, and the user effort required for these two cases is enormous. There should be no need to withstand an off-line attack if the service does an adequate job of protecting the file of hashed passwords [7, 4]. This is an example of creating good alignment between an organizational and user goals [14]; it is known that failure to do so can result in increased non-compliance [9].

On this, and many other questions, we appear to have lost sight of the original goal. We are pursuing substitute goals, such as password complexity and expiration, as ends in themselves long after it has become clear that they do little for the original problem. We persist on a course that was set decades ago, even though the threat landscape has changed beyond recognition and lower-cost alternatives appear available.

Passwords offers egregious examples, but the picture is little better elsewhere. Users are advised to decline certificate warning options independent of any evidence of the relative frequency and costs of true and false positives. It's hard to persuade users that what we ask them to do is not arbitrary and capricious when much of it is, in fact, arbitrary and capricious. It is difficult to see a way forward that restores credibility that does not involve owning up to past mistakes. The fact that we have been so sure, and so wrong, so often suggests that we might profit from asking ourselves "What else have we got wrong?"

3. CONCLUSION

Usable security has many challenges. The techniques and mandates that made sense in a high-assurance world have proved hard to adapt to the vast low-assurance needs of the two billion Internet users. Most by now acknowledge that forcing users to adapt to the technology is not realistic. At the other extreme the hope that "it should just work" is, unfortunately, probably too optimistic. Thus, users will probably need to be engaged in security matters for the foreseeable future.

More is not the answer. It is easy to fall into the trap of thinking that if we find the right words or slogan we can convince people to spend more time on security. Or that usable security offers a bag of tricks to cajole users into increasing effort. We argue that this view is profoundly in error. It presupposes that users are wrong about the cost-benefit tradeoff of security measures, when the bulk of the evidence suggests the opposite. The problem with the product we offer is not simply that it lacks attractive packaging, but that it offers poor return on investment. There are many ways to reduce potential harm with more user effort. Yet, when the answer is always "do more," they don't sound like the response to any question that the user population asks. There is a pressing need however for better protection at the same or lower levels of effort. Rather than techniques to convince users to treat low-value assets as high, we need advice and tools that are appropriate to value.

We suggest that security needs to consider "going green" as far as users are concerned, in the sense of offering only advice that is effort-neutral. If awareness is to be raised then something else, somewhere must be lowered. Measures that demand increases in user time, effort or attention should suggest where the corresponding decreases can be found to balance things out. We need thorough re-evaluation of current practices based on their efficiency in reducing harm, and to reclaim, where possible, the pools of effort that we have wasted on non-productive tasks.

Bio: Cormac Herley is a Principal Researcher at Microsoft research. His current interests include security, usability, economics and data analysis. Herley has a PhD from Columbia University. Contact him at cormac@microsoft.com.

4. REFERENCES

[1] A. Adams and M. A. Sasse. Users are not the Enemy. *Comm. ACM*, 1999.

[2] A. Beautement, M.A. Sasse and M. Wonham. The Compliance Budget: Managing Security Behaviour in Organisations. *NSPW*, 2008.

[3] A. Whitten and J. D. Tygar. Why Johnny Can't Encrypt: A Usability Evaluation of PGP 5.0. *Usenix Security*, 1999.

[4] G. D. Crescenzo, R. J. Lipton, and S. Walfish. Perfectly secure password protocols in the bounded retrieval model. In *Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA,*, pages 225–244, March 2006.

[5] D. Florêncio and C. Herley. Is Everything We Know About Password-stealing Wrong? *IEEE Security & Privacy Magazine. Nov. 2012.*

[6] D. Florêncio and C. Herley. Sex, Lies and Cyber-crime Surveys. *WEIS, 2011, Fairfax.*

[7] D. Florêncio and C. Herley. Where Do Security Policies Come From? *SOUPS 2010, Redmond.*

[8] C. Herley. So Long, And No Thanks for the Externalities: The Rational Rejection of Security Advice by Users. *NSPW 2009, Oxford.*

[9] I. Kirlappos, A. Beautement, and M. A. Sasse. comply or die is dead: Long live security-aware principal agents.

[10] Matt Weir, Sudhir Aggarwal, Michael Collins, Henry Stern. Testing Metrics for Password Creation Policies by Attacking Large Sets of Revealed Passwords. In *Proc CCS*, 2010.

[11] P. Inglesant and M. A. Sasse. The True Cost of Unusable Password Policies: Password use in the Wild. *CHI*, 2010.

[12] R. Anderson and T. Moore. The Economics of Information Security. *Science Magazine*, 2006.

[13] S. Schechter, R. Dhamija, A. Ozment, I. Fischer. The Emperor's New Security Indicators: An evaluation of website authentication and the effect of role playing on usability studies. In *Proc. IEEE Symposium on Security and Privacy*, 2007.

[14] T. S. Teo and W. R. King. Integration between business planning and information systems planning: an evolutionary-contingency perspective. *Journal of management information systems*, 14:185–214, 1997.

[15] Y. Zhang, F. Monrose and M. K. Reiter. The security of modern password expiration: An algorithmic framework and empirical analysis. In *Proc. CCS*, 2010.