

When Does Targeting Make Sense for an Attacker?

Cormac Herley,
Microsoft Research, Redmond

How do so many Internet users escape harm? The range of attacks is enormous and growing; we know that most users neglect even very basic defense measures. Yet things somehow muddle along: two billion people use the Internet and seem to derive more good from it than harm. If security is only as good as the weakest link why don't worst-case outcomes happen regularly? Why isn't everyone hacked every day? The answer may lie in economics rather than technology.

Scalable and Non-scalable attacks

Let's segment attacks into two types, those that scale and those that don't [1]. Scalable attacks have costs that grow much slower than linearly in the number, N , of users attacked. Doubling the number attacked causes the costs to increase by far less than a factor of two: $C_s(2N) \ll 2 C_s(N)$. Thus, the cost of a scalable attack scarcely grows at all with the number attacked. Phishing is scalable, as is any attack that uses spam as the spread vector. Drive-by download attacks, self-replicating viruses and anything that can be completely automated would be scalable, as the cost has very little dependence on the number attacked. Scalable attacks have similar economics to a software product or information good in that first-copy costs dominate [2].

Non-scalable attacks, by contrast, are everything else. Generally they have a linear cost dependence on N . Doubling the number attacked doubles the cost: $C_{ns}(2N) \approx 2 C_{ns}(N)$. Anything that requires per-user effort is non-scalable. Attacks that involve knowledge about the target aren't scalable. For example, the majority of the social engineering attacks described by Mitnick [3] require elaborate target-specific effort. That certainly doesn't scale unless the information can be gathered by a script. Thus, learning the likely answers to backup authentication questions is not scalable. It is far from simple to gather the pet's name, favorite sports team or name of the favorite high-school teacher for a million users in an automated way. Physical side-channel attacks, which require proximity, aren't scalable: getting close to a million people costs a lot more than getting close to one.

This segmentation into scalable and non-scalable attacks is obviously a simplification. Even spam has a linear cost component (e.g. gathering target addresses, finding enough machines and IP addresses to do the sending). However, first-copy costs dominate, so that doubling the

size of the attack has little effect on the overall cost. Equally, attacks that are scalable may need to be followed by a non-scalable component. While phishing may harvest passwords in bulk, the process of cashing out might be non-scalable and one-by-one. Nonetheless, this simplified segmentation of the attack space will prove useful. Interestingly, the attacker investment or team-size to achieve an effect, proves useful in other areas such as attacks on voting systems [8].

Economic Properties

Because of their cost structure, scalable attacks generally reach orders of magnitude more users. A brute-force guessing attack can be directed at the hundreds of millions of accounts that a large web-service might have, while a non-scalable attack, such as shoulder-surfing, might reach only dozens for the same cost. Generic spam campaigns, again, reach hundreds of millions, while carefully researched spear-phishing attacks that are personalized to the recipients might (for the same cost) reach tens. The disparity in terms of reach is enormous. Scalable attacks are non-selective. They attack anyone and everyone. Since the marginal cost per additional user is close to zero, it makes no sense to leave reachable targets un-attacked.

As a consequence, scalable attacks are non-adaptive. Personalization and customization are very limited in a scalable cost model. This is why 419 (Nigerian-style) and lottery scam emails generally begin “Dear Sir/Madam” or “Attn: Beneficiary.” A script can accommodate segments of the population (e.g. a malicious server might attempt different exploits depending on the browser version) but individual-level customization is out.

The objects of scalable attacks are (or become) commodities. First, since they are automated the attack script can be passed to many. As automation improves the skill needed decreases and the pool of potential attackers grows, and so does supply. Thus, the product of scalable attacks often have the race-to-the-bottom economics (where the lowest-price producer dominates) that information goods display [2]. Second, the product of scalable attacks must sometimes be processed by a non-scalable monetization strategy [9]. For example, passwords are harvested in great numbers by scalable attacks, but cashing out involves the (non-scalable) recruitment and management of money mules [6]. If attacks generate product faster than the monetization strategy can process it, a glut ensues and it is to be expected that the price falls. Anecdotal evidence suggests that the prices asked for stolen Credit Card Numbers, passwords, botnet machines and CAPTCHA solving services have indeed been falling. Thus, V_s tends to fall with time. This echoes experience from the regular economy: mass-produced commodities tend to decrease in value over time as the ability to scale up gets better.

Competing against Scalable attacks

Most resources can be attacked in lots of different ways, some of them scalable, and some of them not. Now, if we view attacking as an economic proposition, how do scalable and non-scalable attacks divide the opportunity between them?

Scalable attacks have enviable reach, but offer a very restricted palette of options. Having costs that grow slower than linearly is the exception rather than the rule. The vast majority of attacks do not have this property. On the other hand, there is an almost unlimited suite of non-scalable attacks, but on a per-user basis they are far more expensive. Spam-based attacks, for example, seem to be able to attack users for pennies per million [7]. Non-scalable attacks could easily cost six or seven orders of magnitude more than this per attacked user. Their greater cost suggests that non-scalable attacks must be reserved for cases where the value of the target is extremely high.

Let's be more precise, and examine how a non-scalable attack fares when it competes for the same resource against a scalable one. If N is the number of users attacked, Y is the yield (i.e. fraction who succumb) and V is the average value extracted from those who succumb, then the return, for any attack, is NYV . The yield, average value and number attacked will be different for scalable and non-scalable attacks, so the question is under what circumstances does a non-scalable attack do better? That is when do we have $N_s Y_s V_s < N_{ns} Y_{ns} V_{ns}$.

Now, the non-scalable attacker has a structural disadvantage: he is beaten by orders of magnitude in terms of the number of users he can reach for a given cost, so $N_s \gg N_{ns}$. The only way he can beat the scalable return is if he can make up the difference in yield, or in extracted value, or in some combination of the two.

If the non-scalable attacker extracts the same value per user as his scalable counterpart (i.e., $V_s = V_{ns}$) then he must make up the difference in yield alone. Effectively, then he competes on price: the good produced is the same, so he must achieve lower cost. Recall, however, that scalable attacks produce commodity goods, with race-to-the-bottom economics: V_s tends to fall with time. So, if $V_s = V_{ns}$, both attacks must reduce per-user cost as extracted value falls. Improving automation can deliver these savings for the scalable attacker, but the non-scalable attacker must accept steadily falling return as V_s falls. Thus, while the yield on scalable attacks is often very low (so beating the scalable attacker significantly there is likely) that isn't enough: if $V_s = V_{ns}$ the non-scalable attacker faces constantly decreasing returns.

This suggests that at least some of the orders of magnitude lost in reach will have to be made up in extracted value. So, for successfully attacked users, the extracted value from a non-scalable attack, V_{ns} , must be many times greater than from a scalable one (i.e., $V_s \ll V_{ns}$).

Thus, a non-scalable attack requires two things to compete successfully [1]. First, there must be some users who have much higher extractable value than the average. If the attacker's cost per user is orders of magnitude higher then he needs to extract orders of magnitude more when he succeeds. Second, it must be observable who those high-value users are. It does the attacker no good to know that they exist, if he doesn't know where.

Consider two extreme cases that illustrate the difficulty. First, suppose that value is uniformly distributed: every user has equal value. Non-scalable attacks make no sense in this case since the scalable attacks gather them at far lower cost. At another extreme, suppose that value is concentrated, but entirely unobservable. Again, the non-scalable attacker can't compete if

there are good victims out there, but they just can't be found. Concentration and observability of the extractable value are absolute requirements.

Concentration and Observability of Value

So, when does this happen? When is value concentrated enough that some users have extractable value that is orders of magnitude higher than the average? Worst, as we've seen is if value is uniformly distributed. A little better is any distribution where the variance is high, so that at least some users have significantly higher value. Best are heavytail distributions where a large portion of the overall value lies with a few individuals. Fortunately (for the non-scalable attacker) many phenomena follow power law distributions (e.g. Pareto). Wealth, for example, is power-law distributed, with 1% of the US population owning 35% or so of the wealth.

In these distributions, however, the mean is higher than the median, so most people have below average value. In other words, for the distributions that favor non-scalable attacks, the vast majority of people have below average value. Since the non-scalable attacker needs much higher than average value, he must leave the vast majority of users alone.

When is value observable? Fame is an obvious example. There is little mystery about who is famous. Non-scalable attackers who seek the notoriety of hacking celebrity accounts know exactly where to direct their efforts. Thus, Sarah Palin's email, Lindsay Lohan's Twitter and Scarlett Johansson's phone were each hacked by trawling public information, but the hacks were done at different times by different attackers and involved significant effort. Those attacks weren't scalable.

Wealth is more complicated. It is observable that some are much richer than others, but recall it is extractable wealth that the attacker cares about. Billionaires are easily identified, but the amount extracted isn't necessarily proportional to net worth. As a matter of fact, for consumers, it seems that transactions larger than a few \$100 are more easily detected, and often rolled back [6]. Thus, the distribution is closer to uniform than our attacker would like. Much better targets appear to be small businesses: it is observable that extractable value is concentrated there, and large transactions may arouse less suspicion at a busy company than in consumer accounts.

A very interesting case is where value is concentrated, but not observable. One example is where the value is due to sloppy security practices. For example, many users undoubtedly chose their dog's name as password, or answers for backup authentication questions that are easily learned from available information (e.g. their social network postings). Their lack of care makes them easy targets, but that helps the attacker only if he knows they are easy targets or has already decided to attack them.

Gullibility is another example of a quality that is unobservable. The value for Nigerian, lottery and related scams is concentrated among the most gullible. One can view these scams as

having a scalable spam-campaign front-end which reveals the desired but unobservable quality, by getting the best marks to step forward [4].

Closeness between attacker and victim, and non-economic motives, may increase value. A webmail password might be of little value to a hacker who wishes to monetize it, but be very valuable to a jealous ex-significant other.

Lessons and Conclusion

So what have we learned? The Internet has put billions of people within easy reach of criminals and scammers; that is certainly alarming. However, the Internet has also been cruel to businesses that don't scale. In the legal economy business models that don't scale have been replaced with ones that do, and models that do scale have been replaced with ones that scale better. Personal travel agents gave way to Kayak, Expedia and Travelocity; small booksellers could not compete with Amazon. It's not an accident that it's all but impossible to get a human being on the phone at large web-services: that cost doesn't scale. This trend is unkind to those with linear cost structures. The driving forces in the illegal economy are not different: non-scalable strategies get pushed to the fringes where high margins can be supported.

Only a minority of attacks are scalable, but those that are reach everyone. Most attacks are non-scalable, but the vast majority of users never see them. In this view non-scalable attackers are artisan craftsmen in the age of mass production. When they compete against scalable attacks they must be extremely selective. Unless value is both concentrated and visible their attacks are uneconomic. Even, then only the most valuable targets should be attacked. This suggests a partial answer to puzzle of the missing worst-case outcomes. Users are constantly attacked by the low-yield automated attacks, such as spam and phishing that have become so familiar. However, expensive attacks, such as high-touch social engineering [3], and attacks that require physical proximity [5], are uneconomic at scale and pose little threat to the masses. The 99% or so of users who don't have visibly above average value must guard primarily against scalable attacks. In assessing the economic impact of an exploit or attack we should first ask how well it scales.

[1] C. Herley, The Plight of the Targeted Attacker in a World of Scale, Proc. WEIS 2010.

[2] C. Shapiro and H. Varian, Information Rules, Harvard Business School Press, 1999.

[3] K. Mitnick and W.L. Simon, The Art of Deception, Wiley, 2002.

[4] C. Herley, Why do Nigerian Scammers Say they are from Nigeria? Proc. WEIS, 2012

[5] M. Backes, M. Duermuth and D. Unruh, Compromising Reflections: How to read Computer Monitors around a Corner, IEEE Security & Privacy, Oakland, 2008.

[6] D. Florêncio and C. Herley, Is everything we know about password-stealing wrong? IEEE Security & Privacy Magazine, Dec. 2012.

- [7] C. Kanich, C. Kreibich, K. Levchenko, B. Enright, G. Voelker, V. Paxson, and S. Savage, **Spamalytics: An empirical analysis of spam marketing conversion**, Proc. ACM CCS, 2008.
- [8] E.L. Lazurus, D.L. Dill, J. Epstein and J.L Hall, **Applying a Reusable Election Threat Model at the County Level**, Usenix EVT/WOTE 2011.
- [9] A. Odlyzko, **Providing Security with Insecure Systems**, WiSec 2010.