“The closest the security industry has to a rock star.”
—*The Register*

BRUCE SCHNEIER

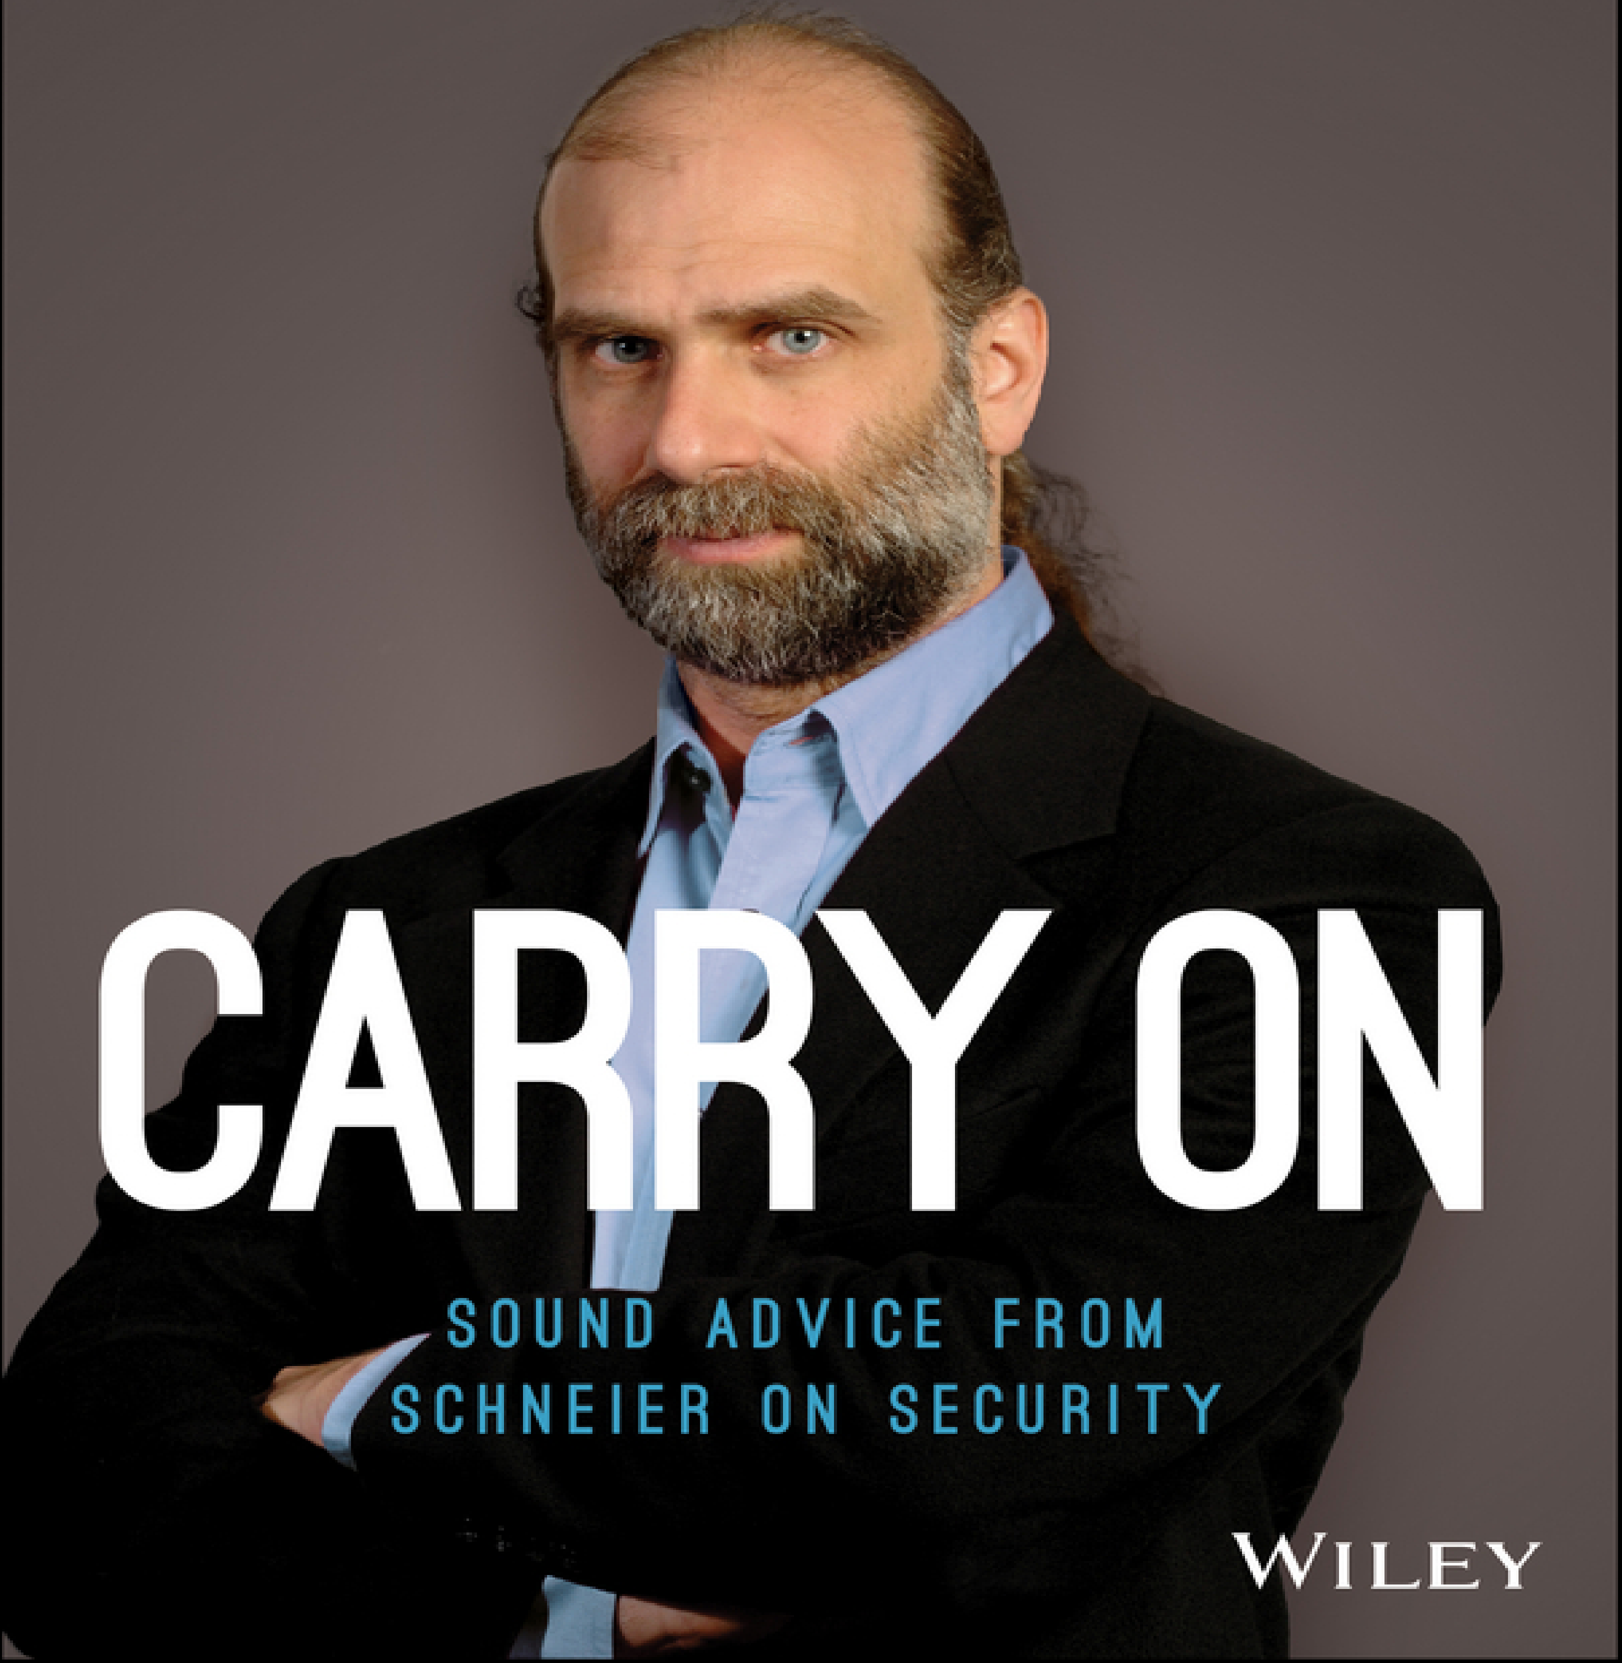

CARRY ON

SOUND ADVICE FROM SCHNEIER ON SECURITY

WILEY

Table of Contents

Chapter 1

The Business and Economics of Security

Consolidation: Plague or Progress

Originally published in Information Security, *March 2008*

This essay appeared as the second half of a point/counterpoint with Marcus Ranum.

We know what we don't like about buying consolidated product suites: one great product and a bunch of mediocre ones. And we know what we don't like about buying best-of-breed: multiple vendors, multiple interfaces, and multiple products that don't work well together. The security industry has gone back and forth between the two, as a new generation of IT security professionals rediscovers the downsides of each solution.

The real problem is that neither solution really works, and we continually fool ourselves into believing whatever we don't have is better than what we have at the time. And the real solution is to buy results, not products.

Honestly, no one wants to buy IT security. People want to buy whatever they want—connectivity, a Web presence, email, networked applications, whatever—and they want it to be secure. That they're forced to spend money on IT security is an artifact of the youth of the computer industry. And sooner or later the need to buy security will disappear.

It will disappear because IT vendors are starting to realize they have to provide security as part of whatever they're selling. It will disappear because organizations are starting to buy services instead of products, and demanding security as part of those services. It will disappear because the security industry will disappear as a consumer category, and will instead market to the IT industry.

The critical driver here is outsourcing. Outsourcing is the ultimate consolidator, because the customer no longer cares about the details. If I buy my network services from a large IT infrastructure company, I don't care if it secures things by installing the hot new intrusion prevention systems, by configuring the routers and servers as to obviate the need for network-based security, or if it uses magic security dust given to it by elven kings. I just want a contract that specifies a level and quality of service, and my vendor can figure it out.

IT is infrastructure. Infrastructure is always outsourced. And the details of how the infrastructure works are left to the companies that provide it.

This is the future of IT, and when that happens we're going to start to see a type of consolidation we haven't seen before. Instead of large security companies gobbling up small security companies, both large and small security companies will be gobbled up by non-security companies. It's already starting to happen. In 2006, IBM bought ISS. The same year BT bought my company, Counterpane, and last year it bought INS. These aren't large security companies buying small security companies; these are non-security companies buying large and small security companies.

If I were Symantec and McAfee, I would be preparing myself for a buyer.

This is good consolidation. Instead of having to choose between a single product suite that isn't very good or a best-of-breed set of products that don't work well together, we can ignore the issue completely. We can just find an infrastructure provider that will figure it out and make it work—who cares how?

Prediction: RSA Conference Will Shrink Like a Punctured Balloon

Originally published in Wired News, *April 17, 2008*

Last week was the RSA Conference, easily the largest information security conference in the world. More than 17,000 people descended on San Francisco's Moscone Center to hear some of the more than 250 talks, attend I-didn't-try-to-count parties, and try to evade over 350 exhibitors vying to sell them stuff.

Talk to the exhibitors, though, and the most common complaint is that the attendees aren't buying.

It's not the quality of the wares. The show floor is filled with new security products, new technologies, and new ideas. Many of these are products that will make the attendees' companies more secure in all sorts of different ways. The problem is that most of the people attending the RSA Conference can't understand what the products do or why they should buy them. So they don't.

I spoke with one person whose trip was paid for by a smallish security firm. He was one of the company's first customers, and the company was proud to parade him in front of the press. I asked him whether he walked through the show floor, looking at the company's competitors to see if there was any benefit to switching.

“I can't figure out what any of those companies do,” he replied.

I believe him. The booths are filled with broad product claims, meaningless security platitudes and unintelligible marketing literature. You could walk into a booth, listen to a five-minute sales pitch by a marketing type, and still not know what the company does. Even seasoned security professionals are confused.

Commerce requires a meeting of the minds between buyer and seller, and it's just not happening. The sellers can't explain what they're selling to the buyers, and the buyers don't buy because they don't understand what the sellers are selling. There's a mismatch between the two; they're so far apart that they're barely speaking the same language.

This is a bad thing in the near term—some good companies will go bankrupt and some good security technologies won't get deployed—but it's a good thing in the long run. It demonstrates that the computer industry is maturing: IT is getting complicated and subtle, and users are starting to treat it like infrastructure.

For a while now I have predicted the death of the security industry. Not the death of information security as a vital requirement, of course, but the death of the end-user security industry that gathers at the RSA Conference. When something becomes infrastructure—power, water, cleaning service, tax preparation—customers care less about details and more about results. Technological innovations become something the infrastructure providers pay attention to, and they package it for their customers.

No one wants to buy security. They want to buy something truly useful—database management systems, Web 2.0 collaboration tools, a company-wide network—and they want it to be secure. They don't want to have to become IT security experts. They don't want to have to go to the RSA Conference. This is the future of IT security.

You can see it in the large IT outsourcing contracts that companies are signing—not security outsourcing contracts, but more general IT contracts that include security. You can see it in the current wave of industry consolidation: not large security companies buying small security companies, but non-security companies buying security companies. And you can see it in the new popularity of software as a service: Customers want solutions; who cares about the details?

Imagine if the inventor of antilock brakes—or any automobile safety or security feature—had to sell them directly to the consumer. It would be an uphill battle convincing the average driver that he needed to buy them; maybe that technology would have succeeded and maybe it wouldn't. But that's not what happens. Antilock brakes, airbags and that annoying sensor that beeps when you're backing up too close to another object are sold to automobile companies, and those companies bundle them together into cars that are sold to consumers. This doesn't mean that automobile safety isn't important, and often these new features are touted by the car manufacturers.

The RSA Conference won't die, of course. Security is too important for that. There will still be new technologies, new products and new startups. But it will become inward-facing, slowly turning into an industry conference. It'll be security companies selling to the companies who sell to corporate and home users—and will no longer be a 17,000-person user conference.

How to Sell Security

Originally published in CIO, *May 26, 2008*

It's a truism in sales that it's easier to sell someone something he wants than a defense against something he wants to avoid. People are reluctant to buy insurance, or home security devices, or computer security anything. It's not they don't ever buy these things, but it's an uphill struggle.

The reason is psychological. And it's the same dynamic when it's a security vendor trying to sell its products or services, a CIO trying to convince senior management to invest in security or a security officer trying to implement a security policy with her company's employees.

It's also true that the better you understand your buyer, the better you can sell.

Why People Are Willing to Take Risks

First, a bit about Prospect Theory, the underlying theory behind the newly popular field of behavioral economics. Prospect Theory was developed by Daniel Kahneman and Amos Tversky in 1979 (Kahneman went on to win a Nobel Prize for this and other similar work) to explain how people make trade-offs that involve risk. Before this work, economists had a model of “economic man,” a rational being who makes trade-offs based on some logical calculation. Kahneman and Tversky showed that real people are far more subtle and ornery.

Here's an experiment that illustrates Prospect Theory. Take a roomful of subjects and divide them into two groups. Ask one group to choose between these two alternatives: a sure gain of $500 and 50 percent chance of gaining $1,000. Ask the other group to choose between these two alternatives: a sure loss of $500 and a 50 percent chance of losing $1,000.

These two trade-offs are very similar, and traditional economics predicts that whether you're contemplating a gain or a loss doesn't make a difference: People make trade-offs based on a straightforward calculation of the relative outcome. Some people prefer sure things and others prefer to take chances. Whether the outcome is a gain or a loss doesn't affect the mathematics and therefore shouldn't affect the results. This is traditional economics, and it's called Utility Theory.

But Kahneman's and Tversky's experiments contradicted Utility Theory. When faced with a gain, about 85 percent of people chose the sure smaller gain over the risky larger gain. But when faced with a loss, about 70 percent chose the risky larger loss over the sure smaller loss.

This experiment, repeated again and again by many researchers, across ages, genders, cultures and even species, rocked economics, yielded the same result. Directly contradicting the traditional idea of “economic man,” Prospect Theory recognizes that people have subjective values for gains and losses. We have evolved a cognitive bias: a pair of heuristics. One, a sure gain is better than a chance at a greater gain, or “A bird in the hand is worth two in the bush.” And two, a sure loss is worse than a chance at a greater loss, or “Run away and live to fight another day.” Of course, these are not rigid rules. Only a fool would take a sure $100 over a 50 percent chance at $1,000,000. But all things being equal, we tend to be risk-averse when it comes to gains and risk-seeking when it comes to losses.

This cognitive bias is so powerful that it can lead to logically inconsistent results. Google the “Asian Disease Experiment” for an almost surreal example. Describing the same policy choice in different ways—either as “200 lives saved out of 600” or “400 lives lost out of 600”—yields wildly different risk reactions.

Evolutionarily, the bias makes sense. It's a better survival strategy to accept small gains rather than risk them for larger ones, and to risk larger losses rather than accept smaller losses. Lions, for example, chase young or wounded wildebeests because the investment needed to kill them is lower. Mature and healthy prey would probably be more nutritious, but there's a risk of missing lunch entirely if it gets away. And a small meal will tide the lion over until another day. Getting through today is more important than the possibility of having food tomorrow. Similarly, it is better to risk a larger loss than to accept a smaller loss. Because animals tend to live on the razor's edge between starvation and reproduction, any loss of food—whether small or large—can be equally bad. Because both can result in death, and the best option is to risk everything for the chance at no loss at all.

How to Sell Security

How does Prospect Theory explain the difficulty of selling the prevention of a security breach? It's a choice between a small sure loss—the cost of the security product—and a large risky loss: for example, the results of an attack on one's network. Of course there's a lot more to the sale. The buyer has to be convinced that the product works, and he has to understand the threats against him and the risk that something bad will happen. But all things being equal, buyers would rather take the chance that the attack won't happen than suffer the sure loss that comes from purchasing the security product.

Security sellers know this, even if they don't understand why, and are continually trying to frame their products in positive results. That's why you see slogans with the basic message, “We take care of security so you can focus on your business,” or carefully crafted ROI models that demonstrate how profitable a security purchase can be. But these never seem to work. Security is fundamentally a negative sell.

One solution is to stoke fear. Fear is a primal emotion, far older than our ability to calculate trade-offs. And when people are truly scared, they're willing to do almost anything to make that feeling go away; lots of other psychological research supports that. Any burglar alarm salesman will tell you that people buy only after they've been robbed, or after one of their neighbors has been robbed. And the fears stoked by 9/11, and the politics surrounding 9/11, have fueled an entire industry devoted to counterterrorism. When emotion takes over like that, people are much less likely to think rationally.

Though effective, fear mongering is not very ethical. The better solution is not to sell security directly, but to include it as part of a more general product or service. Your car comes with safety and security features built in; they're not sold separately. Same with your house. And it should be the same with computers and networks. Vendors need to build security into the products and services that customers actually want. CIOs should include security as an integral part of everything they budget for. Security shouldn't be a separate policy for employees to follow but part of overall IT policy.

Security is inherently about avoiding a negative, so you can never ignore the cognitive bias embedded so deeply in the human brain. But if you understand it, you have a better chance of overcoming it.

Why Do We Accept Signatures by Fax?

Originally published in Wired News, *May 29, 2008*

Aren't fax signatures the weirdest thing? It's trivial to cut and paste—with real scissors and glue—anyone's signature onto a document so that it'll look real when faxed. There is so little security in fax signatures that it's mind-boggling that anyone accepts them.

Yet people do, all the time. I've signed book contracts, credit card authorizations, nondisclosure agreements and all sorts of financial documents—all by fax. I even have a scanned file of my signature on my computer, so I can virtually cut and paste it into documents and fax them directly from my computer without ever having to print them out. What in the world is going on here?

And, more importantly, why are fax signatures still being used after years of experience? Why aren't there many stories of signatures forged through the use of fax machines?

The answer comes from looking at fax signatures not as an isolated security measure, but in the context of the larger system. Fax signatures work because signed faxes exist within a broader communications context.

In a 2003 paper, *Economics*, *Psychology*, *and Sociology of Security*, professor Andrew Odlyzko looks at fax signatures and concludes:

> *Although fax signatures have become widespread, their usage is restricted. They are not used for final contracts of substantial value, such as home purchases. That means that the insecurity of fax communications is not easy to exploit for large gain. Additional protection against abuse of fax insecurity is provided by the context in which faxes are used. There are records of phone calls that carry the faxes, paper trails inside enterprises and so on. Furthermore, unexpected large financial transfers trigger scrutiny. As a result, successful frauds are not easy to carry out by purely technical means.*

He's right. Thinking back, there really aren't ways in which a criminal could use a forged document sent by fax to defraud me. I suppose an unscrupulous consulting client could forge my signature on a non-disclosure agreement and then sue me, but that hardly seems worth the effort. And if my broker received a fax document from me authorizing a money transfer to a Nigerian bank account, he would certainly call me before completing it.

Credit card signatures aren't verified in person, either—and I can already buy things over the phone with a credit card—so there are no new risks there, and Visa knows how to monitor transactions for fraud. Lots of companies accept purchase orders via fax, even for large amounts of stuff, but there's a physical audit trail, and the goods are shipped to a physical address—probably one the seller has shipped to before. Signatures are kind of a business lubricant: mostly, they help move things along smoothly.

Except when they don't.

On October 30, 2004, Tristian Wilson was released from a Memphis jail on the authority of a forged fax message. It wasn't even a particularly good forgery. It wasn't on the standard letterhead of the West Memphis Police Department. The name of the policeman who signed the fax was misspelled. And the time stamp on the top of the fax clearly showed that it was sent from a local McDonald's.

The success of this hack has nothing to do with the fact that it was sent over by fax. It worked because the jail had lousy verification procedures. They didn't notice any discrepancies in the fax. They didn't notice the phone number from which the fax was sent. They didn't call and verify that it was official. The jail was accustomed to getting release orders via fax, and just acted on this one without thinking. Would it have been any different had the forged release form been sent by mail or courier?

Yes, fax signatures always exist in context, but sometimes they are the linchpin within that context. If you can mimic enough of the context, or if those on the receiving end become complacent, you can get away with mischief.

Arguably, this is part of the security process. Signatures themselves are poorly defined. Sometimes a document is valid even if not signed: A person with both hands in a cast can still buy a house. Sometimes a document is invalid even if signed: The signer might be drunk, or have a gun pointed at his head. Or he might be a minor. Sometimes a valid signature isn't enough; in the United States there is an entire infrastructure of “notary publics” who officially witness signed documents. When I started filing my tax returns electronically, I had to sign a document stating that I wouldn't be signing my income tax documents. And banks don't even bother verifying signatures on checks less than $30,000; it's cheaper to deal with fraud after the fact than prevent it.

Over the course of centuries, business and legal systems have slowly sorted out what types of additional controls are required around signatures, and in which circumstances.

Those same systems will be able to sort out fax signatures, too, but it'll be slow. And that's where there will be potential problems. Already fax is a declining technology. In a few years it'll be largely obsolete, replaced by PDFs sent over e-mail and other forms of electronic documentation. In the past, we've had time to figure out how to deal with new technologies. Now, by the time we institutionalize these measures, the technologies are likely to be obsolete.

What that means is people are likely to treat fax signatures—or whatever replaces them—exactly the same way as paper signatures. And sometimes that assumption will get them into trouble.

But it won't cause social havoc. Wilson's story is remarkable mostly because it's so exceptional. And even he was rearrested at his home less than a week later. Fax signatures may be new, but fake signatures have always been a possibility. Our legal and business systems need to deal with the underlying problem—false authentication—rather than focus on the technology of the moment. Systems need to defend themselves against the possibility of fake signatures, regardless of how they arrive.

The Pros and Cons of LifeLock

Originally published in Wired News, *June 12, 2008*

LifeLock, one of the companies that offers identity-theft protection in the United States, has been taking quite a beating recently. They're being sued by credit bureaus, competitors and lawyers in several states that are launching class action lawsuits. And the stories in the media… it's like a piranha feeding frenzy.

There are also a lot of errors and misconceptions. With its aggressive advertising campaign and a CEO who publishes his Social Security number and dares people to steal his identity—Todd Davis, 457-55-5462—LifeLock is a company that's easy to hate. But the company's story has some interesting security lessons, and it's worth understanding in some detail.

In December 2003, as part of the Fair and Accurate Credit Transactions Act, or FACTA, credit bureaus were forced to allow you to put a fraud alert on their credit reports, requiring lenders to verify your identity before issuing a credit card in your name. This alert is temporary, and expires after 90 days. Several companies have sprung up—LifeLock, Debix, LoudSiren, TrustedID—that automatically renew these alerts and effectively make them permanent.

This service pisses off the credit bureaus and their financial customers. The reason lenders don't routinely verify your identity before issuing you credit is that it takes time, costs money and is one more hurdle between you and another credit card. (Buy, buy, buy—it's the American way.) So in the eyes of credit bureaus, LifeLock's customers are inferior goods; selling their data isn't as valuable. LifeLock also opts its customers out of pre-approved credit card offers, further making them less valuable in the eyes of credit bureaus.

And, so began a smear campaign on the part of the credit bureaus. You can read their points of view in this *New York Times* article, written by a reporter who didn't do much more than regurgitate their talking points. And the class action lawsuits have piled on, accusing LifeLock of deceptive business practices, fraudulent advertising and so on. The biggest smear is that LifeLock didn't even protect Todd Davis, and that his identity was allegedly stolen.

It wasn't. Someone in Texas used Davis's SSN to get a $500 advance against his paycheck. It worked because the loan operation didn't check with any of the credit bureaus before approving the loan—perfectly reasonable for an amount this small. The payday-loan operation called Davis to collect, and LifeLock cleared up the problem. His credit report remains spotless.

The Experian credit bureau's lawsuit basically claims that fraud alerts are only for people who have been victims of identity theft. This seems spurious; the text of the law states that anyone “who asserts a good faith suspicion that the consumer has been or is about to become a victim of fraud or related crime” can request a fraud alert. It seems to me that includes anybody who has ever received one of those notices about their financial details being lost or stolen, which is everybody.

As to deceptive business practices and fraudulent advertising—those just seem like class action lawyers piling on. LifeLock's aggressive fear-based marketing doesn't seem any worse than a lot of other similar advertising campaigns. My guess is that the class action lawsuits won't go anywhere.

In reality, forcing lenders to verify identity before issuing credit is exactly the sort of thing we need to do to fight identity theft. Basically, there are two ways to deal with identity theft: Make personal information harder to steal, and make stolen personal information harder to use. We all know the former doesn't work, so that leaves the latter. If Congress wanted to solve the problem for real, one of the things it would do is make fraud alerts permanent for everybody. But the credit industry's lobbyists would never allow that.

LifeLock does a bunch of other clever things. They monitor the national address database, and alert you if your address changes. They look for your credit and debit card numbers on hacker and criminal websites and such, and assist you in getting a new number if they see it. They have a million-dollar service guarantee—for complicated legal reasons, they can't call it insurance—to help you recover if your identity is ever stolen.

But even with all of this, I am not a LifeLock customer. At $120 a year, it's just not worth it. You wouldn't know it from the press attention, but dealing with identity theft has become easier and more routine. Sure, it's a pervasive problem. The Federal Trade Commission reported that 8.3 million Americans were identity-theft victims in 2005. But that includes things like someone stealing your credit card and using it, something that rarely costs you any money and that LifeLock doesn't protect against. New account fraud is much less common, affecting 1.8 million Americans per year, or 0.8 percent of the adult population. The FTC hasn't published detailed numbers for 2006 or 2007, but the rate seems to be declining.

New card fraud is also not very damaging. The median amount of fraud the thief commits is $1,350, but you're not liable for that. Some spectacularly horrible identity-theft stories notwithstanding, the financial industry is pretty good at quickly cleaning up the mess. The victim's median out-of-pocket cost for new account fraud is only $40, plus ten hours of grief to clean up the problem. Even assuming your time is worth $100 an hour, LifeLock isn't worth more than $8 a year.

And it's hard to get any data on how effective LifeLock really is. They've been in business three years and have about a million customers, but most of them have joined up in the last year. They've paid out on their service guarantee 113 times, but a lot of those were for things that happened before their customers became customers. (It was easier to pay than argue, I assume.) But they don't know how often the fraud alerts actually catch an identity thief in the act. My guess is that it's less than the 0.8 percent fraud rate above.

LifeLock's business model is based more on the fear of identity theft than the actual risk.

It's pretty ironic of the credit bureaus to attack LifeLock on its marketing practices, since they know all about profiting from the fear of identity theft. FACTA also forced the credit bureaus to give Americans a free credit report once a year upon request. Through deceptive marketing techniques, they've turned this requirement into a multimillion-dollar business.

Get LifeLock if you want, or one of its competitors if you prefer. But remember that you can do most of what these companies do yourself. You can put a fraud alert on your own account, but you have to remember to renew it every three months. You can also put a credit freeze on your account, which is more work for the average consumer but more effective if you're a privacy wonk—and the rules differ by state. And maybe someday Congress will do the right thing and put LifeLock out of business by forcing lenders to verify identity every time they issue credit in someone's name.

The Problem Is Information Insecurity

Originally published in Security Watch, *August 10, 2008*

Information insecurity is costing us billions. We pay for it in theft: information theft, financial theft. We pay for it in productivity loss, both when networks stop working and in the dozens of minor security inconveniences we all have to endure. We pay for it when we have to buy security products and services to reduce those other two losses. We pay for security, year after year.

The problem is that all the money we spend isn't fixing the problem. We're paying, but we still end up with insecurities.

The problem is insecure software. It's bad design, poorly implemented features, inadequate testing and security vulnerabilities from software bugs. The money we spend on security is to deal with the effects of insecure software.

And that's the problem. We're not paying to improve the security of the underlying software. We're paying to deal with the problem rather than to fix it.

The only way to fix this problem is for vendors to fix their software, and they won't do it until it's in their financial best interests to do so.

Today, the costs of insecure software aren't borne by the vendors that produce the software. In economics, this is known as an externality, the cost of a decision that's borne by people other than those making the decision.

There are no real consequences to the vendors for having bad security or low-quality software. Even worse, the marketplace often rewards low quality. More precisely, it rewards additional features and timely release dates, even if they come at the expense of quality.

If we expect software vendors to reduce features, lengthen development cycles and invest in secure software development processes, it needs to be in their financial best interests to do so. If we expect corporations to spend significant resources on their own network security—especially the security of their customers—it also needs to be in their financial best interests.

Liability law is a way to make it in those organizations' best interests. Raising the risk of liability raises the costs of doing it wrong and therefore increases the amount of money a CEO is willing to spend to do it right. Security is risk management; liability fiddles with the risk equation.

Basically, we have to tweak the risk equation so the CEO cares about actually fixing the problem, and putting pressure on his balance sheet is the best way to do that.

Clearly, this isn't all or nothing. There are many parties involved in a typical software attack. There's the company that sold the software with the vulnerability in the first place. There's the person who wrote the attack tool. There's the attacker himself, who used the tool to break into a network.

There's the owner of the network, who was entrusted with defending that network. One hundred percent of the liability shouldn't fall on the shoulders of the software vendor, just as 100% shouldn't fall on the attacker or the network owner. But today, 100% of the cost falls directly on the network owner, and that just has to stop.

We will always pay for security. If software vendors have liability costs, they'll pass those on to us. It might not be cheaper than what we're paying today. But as long as we're going to pay, we might as well pay to fix the problem. Forcing the software vendor to pay to fix the problem and then pass those costs on to us means that the problem might actually get fixed.

Liability changes everything. Currently, there is no reason for a software company not to offer feature after feature after feature. Liability forces software companies to think twice before changing something. Liability forces companies to protect the data they're entrusted with. Liability means that those in the best position to fix the problem are actually responsible for the problem.

Information security isn't a technological problem. It's an economics problem. And the way to improve information technology is to fix the economics problem. Do that, and everything else will follow.

Security ROI: Fact or Fiction?

Originally published in CSO Magazine, *September 2, 2008*

Return on investment, or ROI, is a big deal in business. Any business venture needs to demonstrate a positive return on investment, and a good one at that, in order to be viable.

It's become a big deal in IT security, too. Many corporate customers are demanding ROI models to demonstrate that a particular security investment pays off. And in response, vendors are providing ROI models that demonstrate how their particular security solution provides the best return on investment.

It's a good idea in theory, but it's mostly bunk in practice.

Before I get into the details, there's one point I have to make. "ROI" as used in a security context is inaccurate. Security is not an investment that provides a return, like a new factory or a financial instrument. It's an expense that, hopefully, pays for itself in cost savings. Security is about loss prevention, not about earnings. The term just doesn't make sense in this context.

But as anyone who has lived through a company's vicious end-of-year budget-slashing exercises knows, when you're trying to make your numbers, cutting costs is the same as increasing revenues. So while security can't produce ROI, loss prevention most certainly affects a company's bottom line.

And a company should implement only security countermeasures that affect its bottom line positively. It shouldn't spend more on a security problem than the problem is worth. Conversely, it shouldn't ignore problems that are costing it money when there are cheaper mitigation alternatives. A smart company needs to approach security as it would any other business decision: costs versus benefits.

The classic methodology is called annualized loss expectancy (ALE), and it's straightforward. Calculate the cost of a security incident in both tangibles like time and money, and intangibles like reputation and competitive advantage. Multiply that by the chance the incident will occur in a year. That tells you how much you should spend to mitigate the risk. So, for example, if your store has a 10 percent chance of getting robbed and the cost of being robbed is $10,000, then you should spend $1,000 a year on security. Spend more than that, and you're wasting money. Spend less than that, and you're also wasting money.

Of course, that $1,000 has to reduce the chance of being robbed to zero in order to be cost-effective. If a security measure cuts the chance of robbery by 40 percent—to 6 percent a year—then you should spend no more than $400 on it. If another security measure reduces it by 80 percent, it's worth $800. And if two security measures both reduce the chance of being robbed by 50 percent and one costs $300 and the other $700, the first one is worth it and the second isn't.

The Data Imperative

The key to making this work is good data; the term of art is “actuarial tail.” If you're doing an ALE analysis of a security camera at a convenience store, you need to know the crime rate in the store's neighborhood and maybe have some idea of how much cameras improve the odds of convincing criminals to rob another store instead. You need to know how much a robbery costs: in merchandise, in time and annoyance, in lost sales due to spooked patrons, in employee morale. You need to know how much not having the cameras costs in terms of employee morale; maybe you're having trouble hiring salespeople to work the night shift. With all that data, you can figure out if the cost of the camera is cheaper than the loss of revenue if you close the store at night—assuming that the closed store won't get robbed as well. And then you can decide whether to install one.

Cybersecurity is considerably harder, because there just isn't enough good data. There aren't good crime rates for cyberspace, and we have a lot less data about how individual security countermeasures—or specific configurations of countermeasures—mitigate those risks. We don't even have data on incident costs.

One problem is that the threat moves too quickly. The characteristics of the things we're trying to prevent change so quickly that we can't accumulate data fast enough. By the time we get some data, there's a new threat model for which we don't have enough data. So we can't create ALE models.

But there's another problem, and it's that the math quickly falls apart when it comes to rare and expensive events. Imagine you calculate the cost—reputational costs, loss of customers, etc.—of having your company's name in the newspaper after an embarrassing cybersecurity event to be $20 million. Also assume that the odds are 1 in 10,000 of