

Free-Market Security
Leveraging Economics to Improve Software Security and Reduce Risk

Robert J. Brown
rjb@robertjbrown.com
February 24, 2010

As I write this, there is a massive recall and public outcry against Toyota for a faulty accelerator that could cause unintended acceleration. This presents a risk of accident or death in a number of cases and has been taken very seriously by the government, public, and media. My first reaction was this: they should put their CIO/CISO in charge of the recall because they deal with ‘recalls’ multiple times a week in the form of vulnerable software.

Are software vulnerabilities any less risky than a faulty accelerator? Does software not control every major facet of our critical infrastructure, transportation, financial, and personal health and well being? Imagine the highway was filled with cars that have the same number of ‘severity 5’ defects that our software and applications have. How safe would you feel driving home? Would you be willing to take your car in monthly on “Recall Tuesday” to have it fixed?¹

If we have established that software defects and vulnerabilities (which could be misconfigurations, programming errors, and the like) are critical to our well being and economic viability, why do we continue to make choices to purchase new software and develop new applications that are not secured to the level of risk we wish to accept? It seems that we would want to consider security and reliability as one of the cornerstones of our decision-making process, yet we rarely do.

My personal conclusion to that question is because we have the economics wrong. The risk reduction incentives of safer software aren’t aligned with the business decisions when choices are being made. This includes choice of what vendor to work with, what software to purchase, how to develop your own application, how to configure your server, and all of the other factors that contribute to our technical vulnerabilities.

It doesn’t have to be this way. There are models that have been effective in realigning choice and incentives to achieve a goal. Let’s take one specific example as a case study in redefining the incentives to realize a desired outcome.

The Kyoto Protocol – Market-based Pollution Reduction

¹ This is not meant in any way as a slight to Microsoft. They have done a great job with improving security and reducing risk. This was just an easy example of a date/process that readers would be aware of.

The Kyoto Protocol is a treaty established on an international scale and enforced as of 2005. This legally-binding agreement requires the current 183 participant countries² to a 'cap and trade' system to combat the emission of greenhouse gases. The Kyoto Protocol leverages economic incentives to control pollution on a global scale. To summarize:

- The scope of what pollutants to reduce was established to include four greenhouse gases³ and two groups of gases⁴.
- The scope of what countries should participate in the reduction was established based on the United Nations Framework Convention on Climate Change. This generally included the industrialized nations and was expanded to include the current 183 participants.
- The measurement of the baseline level of pollution was taken based on emission information from 1990. This measurement is calculated based on each individual country and their emission level for the gases.
- The amount of greenhouse gas reduction against the 1990 baseline was set to 5.2% for each country measured over the years 2008-2012. Accounting for continued economic and pollutant growth as compared to expected levels in 2010 without the treaty, this represents a 29% reduction from the 1990 level.
- A uniform measurement standard was agreed upon and centralized regulators were created to accept the ongoing measurements provided by each country. The organizations are designed to only oversee the measurement and reporting standards and do not participate or take political positions on climate issues.
- Participating countries are able to reduce their meet their agreed upon limitations by directly reducing emissions, purchasing emission reductions credits from elsewhere, or by initiating projects that reduce emissions in non-Annex I participants.
- In Europe, the European Union Emission Trading Scheme was established. Under the EU ETS, large emitters of gases within the EU must monitor and annually report their CO₂ emissions, and they are obliged every year to return an amount of emission allowances to the government that is equivalent to their CO₂ emissions in that year.
- There are five international exchanges allowing for the sale and purchase of emission credits - Chicago Climate Exchange, European Climate Exchange, Nord Pool, PowerNext and the European Energy Exchange.

The net result of the described Kyoto protocol is to create economic *incentives* for countries and their industrialized companies to reduce emissions. This is achieved without specific global regulations which may help or harm some of the participants, but rather by taking a free market approach to the problem. Polluters are still able to pollute, however it may become more expensive over time if carbon credits are necessary. Buyers of the credits are penalized for polluting while sellers are rewarded for having reduced emissions. The impact is to influence current and future behavior and choices to align them with the overall goal of reducing pollution.

² The United States of America is currently not one of the 183 participants.

³ Carbon dioxide, methane, nitrous oxide, sulphur hexafluoride

⁴ Hydrofluorocarbons and perfluorocarbons

The New Protocol

How could the Kyoto Protocol and cap-and-trade relate to information security? It may take a leap of thinking but here's the idea:

I propose that software vulnerabilities are similar in nature to pollutants – unwanted side effects of bad choices that threaten our environment.

In the same manner that industrial companies wish to quickly launch their factory to begin recognizing revenue at the cost of polluting the environment, the goal of many software development efforts is to release the product or application to market and begin to recognize value – be it revenue or business process enhancements. In many cases, security is an afterthought in this process and something that is deemed to be revisited in a later release or isn't properly considered up-front.

Despite our best efforts to embed security into the planning and quality assurance process through initiatives like threat modeling, in many organizations we have failed to achieve true secure development practices due to competitive and market pressures. The same is true of our server and workstation configurations. To the recipient of bad software design or implementation, it results in a permanent state of identifying vulnerabilities, quantifying them in a rational manner, fixing/patching them, or accepting the risk of not fixing the issues. Information Security professionals are consistently fighting this losing battle.

Perhaps it doesn't need to be that way. With a bit of smart regulation and cooperation, an information security policy and standard could be adopted to change behaviors and align them towards vulnerability management. Using the Kyoto framework as described above, here's how it might work for information security:

- The scope of what vulnerabilities to reduce would be established. This could be defined by anything from severity to industry type. They could be class-based according to CVE identifiers as published by Homeland Security.
- The scope of what countries should participate would be established in a similar collaborative manner to Kyoto. I would propose all first-world economies and major contributors to critical infrastructure components across the world.
- The measurement of the baseline level of vulnerability could be established based again on CVE or criticality for the most recent year. Numerous security vendors also have significant data that could be used as input to this process (from McAfee/Symantec to Google and their Safe Search.)
- The amount of vulnerability reduction against the baseline would be set to a specific percentage for each country/industry and measured over a period of time.
- A uniform measurement standard would be agreed upon and centralized regulators/service providers would be established or engaged to perform or accept the measurements provided by each country. These organizations would be

designed only to oversee the measurement and reporting standards and would not otherwise participate or take political positions on security issues.

- Participating countries would be able to reduce their meet their agreed upon limitations by directly reducing vulnerabilities, purchasing vulnerability reduction credits from elsewhere, or by initiating projects that reduce vulnerabilities in participating locales.

Similar to Kyoto, there are no global regulations on vulnerabilities that could help or harm organizations. Everyone is still able to deploy vulnerable software or applications, however those choices may become more expensive over time if it requires going to the market to purchase credits. Buyers are penalized for their aggregate vulnerabilities while sellers are rewarded for their improved security posture and lower risk profile. The impact is to influence current and future behavior and choices to align them with the overall goal of reducing pollution.

What might happen with a system like this? Companies would start taking a much closer look at their software/application/development/implementation choices. Vendors would start to align against a standard and quantify their vulnerability metrics and security of their offerings. Finally, the global number of vulnerabilities would likely start to decline (factoring in technology growth.) This solution is agnostic to internally-developed applications vs. purchased applications and it works locally or in a cloud. It just means that whatever new technology you come up with, it is evaluated against the baseline criteria and recorded as part of your overall vulnerability total. Measurement becomes a key metric with negative economic impacts being the driver for behavioral change.

Before the arguments start, and there are many, I acknowledge this won't solve every security problem, it won't be 100% accurate, it won't discover every vulnerability, and it does not take into account that a severity 5 vulnerability on a test system may not present the same level of risk as one on a critical production system. It also may not be fair and equitable to a larger organization that has a harder time with security and vulnerability management than a smaller company. I agree with those assertions, but that's not the point. The goal is realignment of business decisions to risk with a focus on vulnerability reduction. Just like the reduction in pollution from climate change, a reduction in vulnerabilities means we have reduced, but not eliminated, risk.

The key concept underlying this idea also relies upon the belief that organizations would report in aggregate their number of vulnerabilities to an outside source. But let's be realistic: there is no organization today that doesn't have vulnerabilities in their environment. Hiding aggregate security information doesn't help to solve the problem and it doesn't make you any more or less vulnerable to attack or compromise.

Would this work to move security and vulnerabilities to the forefront of the business decision-making process? I think it would. How difficult would it be to implement? It would take a lot of coordination, but it could happen with the right level of support from business and government alike. Regulations and international treaties would be an ideal approach to implementing an idea such as this.

In conclusion, this is simply a starting point for a discussion about economics, choice, and its impact on our ability to manage vulnerability and risk. I'm now going to walk outside, get into my Toyota, and hopefully make it home safely.