COMMENTS ON THE NIST PRELIMINARY CYBERSECURITY FRAMEWORK

To Whom It May Concern:

First, thank you NIST, DHS, and NSS staff who contributed to what has been, in my opinion, a very successfully run cybersecurity partnership endeavor. I have attended all of the sessions, have appreciated the dialogue both online and off, and have learned quite a bit. And also, thank you again for taking the time to read these final comments at the end of what must have seemed a long process.

With regard to this document, I have elected to break my response down into four sections: A problem space assessment, discussion of whether the framework as it is is the tool we needed to build, comments on the quality of the tool we did build, and externality/ecosystem thoughts. I will attempt to answer the direct questions asked in these sections - if indirectly.

PROBLEM SPACE

One of the gaps in the framework process, in my opinion, has been in the determination and socialization of a problem definition - I feel like that was not done proactively. Building a "framework" means many things to many people. "Making things more secure" is even more nebulous - even within the context of Executive Order definitions and scoping. I believe many of my concerns and criticisms could have been alleviated by the government spending more time and effort educating participants on the hows and why's and what's of the Executive Order and the Framework up front. Not everyone had a strong historical background and it was evident at the workshops.

That said, rather than making that same mistake here and providing comments without context, I have included several assertions below that have guided my thinking and evaluation of where we are with the framework. Your mileage may vary:

1. We are failing at cybersecurity. If we were not, the topic would not be receiving this much attention or funding and we would not be going through this process. As a nation, just looking at the number of substantial public breaches makes this very clear.
2. We are heavily investing in cybersecurity. Many organizations spend vast sums and use known best practices regularly in their environments.
3. These organizations are still getting breached at unacceptable (to the nation and to themselves) rates despite their security efforts
4. Asking for their "best practices" and "what works for them" typically results in a list of things that, by themselves, are not effectively protecting us from cybersecurity risks
5. Very few, if any, organizations are "doing security" using a substantially different model than anyone else. For the most part, differences are a matter of degree, not philosophical approach
6. Given the high-end investment with lack of consistent success in reducing risk to acceptable levels in those quarters, it is safe to assume that merely eliciting "best practices" and re-framing them will lead to a continued lack of success.
7. In many cases, it appears that this failure occurs in several areas: a) It's very difficult to follow these best practices consistently over time. b)Existing practices assume a level of non-security business, IT, and operational maturity that is not addressed by most best practices and is certainly not often culturally socialized, and c)Expecting thousands of organizations to be able to fund or maintain the level of operational and technical maturity required to substantially reduce risk the way we currently build/implement IT/ICS systems is a potentially irrational expectation - even when everyone is on board with the idea.

8. Despite continued failures, it is still very difficult to elicit a comprehensive view of cybersecurity that does not focus on things like incident response, firewalls, or other "best practice controls" found in documents like 800-53.
9. Open ended elicitation attempts, lacking guidance, typically result in the re-documentation of this "tribal" knowledge and cover no new ground. Cybersecurity is, and continues to be, a discipline that is poorly understood or defined outside of its historically technical focus. This is gap that must be remediated or all future efforts will ultimately fail.
10. Starting with and then breaking down the problem into its fundamental business and cultural considerations - not controls - is critical to identifying and remediating these choke points. Ie, instead of focusing on the execution of security, successful efforts will improve the "environment in which cybersecurity occurs".
11. An example of a method more likely to elicit the environmental/cultural/business underpinnings of cybersecurity would be to start with business outcome (enablement) objectives - any of them will do since security controls tend to collapse/overlap across objectives -- and then elicit a framework to meet these objectives while assuming a lack of a dedicated security team and without making references to cybersecurity specific technologies. This forces a dialogue shift onto new ground and the resulting work could then be linked to existing "hard" security.
12. There is a difference between "improving the strategic relationship between everyone and the bad guys" and "defending my organization today". This difference must be consciously and explicitly articulated and enforced - the framework and EO, in my opinion, are aimed at the former while many practitioners are so busy fighting fires that they're focused on the latter. This shows in resulting work.
13. The length of time it takes to gather consensus opinion on things like the framework - which is filled with content that for the most part already existed - supports the view that there is an underlying reference model needed to get all parties on the same page that has not been articulated, documented, or socialized.
14. Cybersecurity is ultimately a quality assurance problem that consists of a rate (flaws over time). Rate problems cannot all be solved by doing more of the same. I cannot pay someone enough money to, say, flap their arms and fly. The most helpful frameworks, instead of focusing on how we are flapping our arms, actually spend time to break down the problem into it's most basic pieces and help us solve for flight in an engineered, not ad hoc, way.
15. Until we've solved the quality problem, risk management is a red herring. Risk management allows us to identify risks so that we can pivot towards protecting ourselves against those threats we care about most. But, if our organizations are too immature to pivot - have too low a quality of security - the prioritization process will not help us.

IS THE FRAMEWORK THE TOOL WE NEEDED?

Based on the above assertions, it should be clear that I think not. Originally, I was very hopeful. The executive order spelled out the need for outcome objectives which could be met by using the framework. There was talk of linking these objectives - potentially derived from approaches like DHS's CARMA - to business maturity domains (like those found in C2M2.) This was exciting and would have provided an approach to deriving a critical knowledge intersection that is so often lacking in large scale security efforts: non-security specific business needs with non-security specific business structure. This intersection is what helps clearly communicate security to executive stakeholders while allowing those stakeholders to provide the rest of the business with clear success metrics and implementation rails.

Instead, the framework does not help define outcomes, it supports technologists while making vague references to business, it has no methodology or philosophy that speaks to cost-effectiveness, and it completely fails to describe a functional link between cyber and business risk (although I do not believe the two should be different). The language is written in the same terms as other control catalogues that business leaders do not understand (raising the level of abstraction is not the same as putting it in business terms), and the structure of the document reflects a cybersecurity practitioner's view of incident response, not the structure of a business. When it comes time to implement the framework, it will suffer from exactly the same difficulties as other control catalogues. The control language is vague enough that it will be difficult to use it as a "Rosetta Stone", although it could be useful as a link between other frameworks for contract purposes.

We needed a strategic, business-centric framework to help conceptualize the different aspects of security that were not being addressed by other security-specific documents. We needed to structure cybersecurity conceptually so that everyone can get out of each other's way, normalize the conversations, and solve for reduced risk. Instead, we ended with - in my opinion - a rehash of existing material that gives, at best, lip service to some of the other higher level considerations.

Ultimately. I believe the framework does little to change business as usual. For the next round, I believe we should look at this version of the framework, turn it inside out, and include everything except what is in this version.

THE TOOL WE BUILT

Despite the nature of the framework being a miss on what I think we needed, it at least doesn't attempt to go very far beyond being a control catalogue (despite language suggesting it does.) This allows us to take the base and create products on top of it that can provide different, more targeted views - such as a re-write for small utilities. I hope there will be a space created for framework users to store and share these "views" or "lenses" that they create.

One substantial structural problem, though, is the incident response functional layout. There are many reasons why I believe the functions will make consistent implementation difficult, but that would require a longer document - so I'll leave this suggestion: Re-frame it around organizational structures and roles. This will allow different people and parts of an organization know what they need to be doing, how they need to do it, and how they fit into the larger picture. This is a much more business friendly approach and will actually improve efficacy of implementations. It also creates a better metrics environment by allowing the treatment of role/responsibility failures as vulnerabilities which can be managed by the C-suite.

Another suggestion is to, in the future, take a modular policy management approach by letting the framework speak at a very high level and then create more specific technical/process specs which can be updated independently of and much more frequently than the framework itself. This shouldn't be too much of a stretch from how it's written now. This will help solve the "it's too high level!" "it's too specific!" problem and make it easier to maintain over time.

EXTERNALITY/ECOSYSTEM

While my above comments may be read as exceedingly critical, I still view the entire process as a success. As first round efforts go, the government and industry came together in what is, by typical measures a rush and created a viable first round product. This initiative, the resulting momentum,

and the earned goodwill and trust are perhaps more key to the reduction of cybersecurity risk as any framework that could have been created. We have done good work. Further, it would be naive to believe that a first round effort such as this would product a document too far off from center and I am looking forward to the next round.

On that note, and although most of these comments have been written from my own, individual, perspective, we (Energysec) as a non-profit community in the electric sector, believe strongly in this process, the framework, and the ecosystem going forward. In particular, we believe non-profits may serve in key roles. We are interested in providing input into how this framework gets implemented and how to manage/support the required trusted communities that will be involved.

Thank you again for your time, consideration, knowledge, and patience.

Jack Whitsitt | Energysec | jack@energysec.org | |http://twitter.com/sintixerr | 12/13/2013