Feedback on the NIST Cybersecurity Framework 2.0 Initial Public Draft

Raymond Sheh[1] and Karen Geappen[2]

2023-11-06

Introduction

We applaud the tremendous work that the U.S. National Institute of Technology (NIST) and collaborators have put into the new Cybersecurity Framework (CSF) 2.0, and its precursor and associated documents. We can recognise the effort taken to increase the ability to integrate the CSF 2.0 with other international standards and frameworks and in sectors other than Critical Infrastructure. It is heartening to think that these resources will be built on by organizations in both Critical Infrastructure and other sectors all over the world, over the next decade. The CSF 2.0 has the potential to induce a step-function improvement in cybersecurity, right at the moment when malicious cyber actors are beginning to threaten the everyday life of modern society on a large scale.

We recognize that a framework must be concise, domain and technology agnostic, and be readable by a wide audience. We also understand that there are also precious few opportunities to make stakeholders and decision makers aware of risks that they may be otherwise unknowingly taking. It is through this lens that we would like to make some suggestions for further improvement in the Initial Public Draft, in the hope that it can better highlight some real, emerging, cybersecurity risks while still being a concise and general document.

Our feedback[3] is divided into six categories, summarized below and expanded upon in the rest of this document. These categories are:

- Cybersecurity is everyone's problem.

  All too often we hear people and organizations say that they are not likely to be a target or that they don't have information of value. We feel that the CSF 2.0 presents an opportunity to encourage all people to recognize that they are stakeholders in the process and that the risks are interconnected.

[1] Raymond Sheh is the Uncrewed Aircraft Systems Research Lead at the Public Safety Communications Research Division of NIST, and an Adjunct Associate Research Scientist at the Whiting School of Engineering at Johns Hopkins University. He can be reached at [redacted]
[2] Karen Geappen is a Director at Anchoram Consulting, Western Australia. She can be reached at [redacted]
[3] The authors provide this feedback in their personal capacity. The opinions and views expressed in this document do not reflect those of any organization or employer, past or present.

- Stronger organizational context and guardrails.

  Current technology and supply chain practices mean organizations must use systems and processes that are less than ideal from a cybersecurity risk perspective. This may be because these systems are ‘black boxes’, Artificial Intelligence (AI) or otherwise, or because the organization does not have direct control over them. We feel that the CSF 2.0 can benefit from covering guardrails to explicitly recognise, document and communicate clear boundaries between the acceptable and not acceptable in relation to behaviors, outcomes and impacts. We feel that there must be guidance from the CSF 2.0 assisting organizations in recognising, assessing and mitigating risks associated with the loss of visibility within the ‘black box’ or loss of direct control, where discontinuing use is not a viable option.

- Extending the data lifecycle.

  The CSF 2.0 appears to be mostly concerned with managing risks associated with transmitting, using, and disposing of data. We feel that this should be extended to the origination of the data, particularly in the current climate where data is too plentiful for viable human oversight, and used in ways that may not be well understood.

- Cybersecurity and AI (and other) risks are greater than the sum of their parts.

  Cybersecurity risks cannot be considered in isolation, for complexities in business processes can interact with cybersecurity risk in ways that result in risks that are greater than the sum of their parts. In this section we use the interaction between Cybersecurity and AI risk as an illustrative example, to discuss risks across the organization that may not be obvious if only considered from a Cybersecurity or AI risk management perspective alone.

- Cybersecurity has a human factor.

  Cybersecurity starts and ends with humans. We feel that the CSF 2.0 can benefit from a more explicit acknowledgement of individual and organizational human failings that go beyond simply not achieving a particular tier, for instance.

- There can be too much cybersecurity.

  There is no such thing as perfect security. Rather, there is an “appropriate” level of cybersecurity for a given organization. The CSF 2.0 quite rightly focuses on the case where an organization needs to reduce cybersecurity risk. As implied by the call for balance among broader risk management, there is also the possibility that an organization is too cyber risk averse, at the expense of other risks. We feel that this possibility, and the role of the CSF 2.0 in helping an organization to strike a better overall balance that may involve an increase in an organization’s cyber risk, while keeping within appropriate levels, could be acknowledged more clearly.

Note that while we have included more specific examples and commentary for the purpose of illustration, we also realize that much of the detail is outside the scope of the CSF 2.0 and belongs in other associated resources. We include this detail to help the reader to better understand the context within which our suggestions exist.

Cybersecurity is everyone's problem

One argument that we often hear when we discuss cybersecurity with non-cybersecurity folks is that "I'm not a target, no-one cares about me and if they got in, there's nothing important anyway.". Worryingly, we hear about this not just from private individuals, but also from sectors that do have critical or highly personal information, such as emergency response, childcare/childhood education, or community groups.

Usually such people don't seem to realize that they are connected to more than just themselves and their immediate part of the supply chain and that they pose risks to others just as others pose risks to them. For example, it often doesn't occur to such people that a malicious actor could compromise their device and then use it to launch an attack on another entity, perhaps one that they are associated with or perhaps one that is completely unrelated. This presents a risk to both society, as well as themselves as the finger of blame may now be pointed at them. They may also not fully understand how important they may be to other stakeholders and not fully appreciate the effect that an attack on them may pose.

We feel that good opportunities to make this point may be in Sections 1.1 (Audience) and 3 (Using the Framework). The commentary around Table 3, particularly with regard to "Third-Party Cybersecurity Risks", may also benefit from an emphasis that the "third party", who poses risk to them, and who they may pose risks to, may, potentially, have no business relationship at all with the user's organization.

Stronger organizational context and guardrails

Within Appendix C, Table 5, the description of the Govern function, Organizational Context category (GV.OC), includes "The circumstances - mission, stakeholder expectations, and legal, regulatory, and contractual requirements - surrounding the organization's cybersecurity risk management decisions are understood". We feel that this description, and the subcategories, focus on what "should" happen, but perhaps do not stress highly enough the need to determine what "must not" happen. This is implied, for instance through considering requirements and contracts, but is not specifically called out.

We feel that explicitly considering the guardrails that an organization's processes should never breach is a vital part of informing all of the other aspects of the CSF 2.0 functions. The thought processes required to establish these, and the documentation of them, assist in other functions of the CSF 2.0. They can also give rise to improved transparency including for Third Party risks.

Similarly we feel that in Table 7, the Protect function, Platform Security (PR.PS) and Technology Infrastructure Resilience (PR.IR) categories should also have subcategories that explicitly address the need for guardrails to deal with business realities. It may not be possible to maintain, replace, or remove software commensurate with risk if the software is business critical and is, for instance, a legacy system for which the vendor no longer exists and where it is not economically viable to replace it.

This can also occur for an AI system whose behavior cannot be rigorously checked and is therefore a “black box” to traditional risk assessment. This potentially means that unacceptable risks cannot be identified. Similarly, some networks and environments might, by necessity, be more open to potential unauthorized access or threats as part of their business purpose, such as ones that interface with experimental systems.

This can also occur with the increased use of outsourcing, third parties and cloud where direct system control has been handed to others through contracts and agreements. This inherently introduces risks related to trust in the other entity and interpretation of nuances critical for the assessment of risk and incidents. While contracts, regulations, and requirements should theoretically ensure that these risks are managed, ultimately their writing, and interpretation, are imperfect and pose a residual risk. Policy and technical guardrails are an important part of managing this residual risk.

Without explicitly accounting for these business realities, there is a real risk that the user of the CSF 2.0 may miscategorize or ignore particular threats, simply consider them “commensurate with risk” without enacting any policies at all, or may enact policies that are impossible to follow and thus are also ignored. Instead, it may be more useful to explicitly acknowledge that it may be necessary to place guardrails around their use and behavior, sandbox them, monitor them more closely, or otherwise manage the risk. Explicit action to recognise and document things that ‘must not’ or ‘can not’ occur gives rise to awareness and choice on subsequent actions in other CSF 2.0 functions.

For example, an additional subcategory under PR.PS might be “Policies, guardrails, and technical controls are applied to systems that are business critical but by themselves pose an unacceptable risk.”. Similarly, an additional subcategory under PR.IR might be “Networks, environments, and assets that may be exposed to additional risk as part of unavoidable business requirements are accurately identified, characterized, and have policies, guardrails, and technical controls applied to manage risk.”.

Extending the data lifecycle

Also within Appendix C, Table 7 describes the Protect function, Data Security category (PR.DS), which includes “Data is managed consistent with the organization’s risk strategy to protect the confidentiality, integrity, and availability of information”. We feel that the subcategories as written focus on the storage, transmission, use, destruction, and availability of the data, but neglect to also consider cybersecurity risks associated with the origination of the data, be it from within or

outside the organization. This implies recognition that there needs to be trust in the truth of the data - especially when it is used within processes or for decision making. This is different to integrity, which is trust that the data has not been manipulated inappropriately.

Some would argue that the authenticity of incoming data is more of a business risk than a cybersecurity risk. However we feel that there are enough ways in which a malicious actor may create a cybersecurity impact through manipulation of incoming data that it deserves a mention in this context. These risks go beyond the traditional ways that a malicious actor may use compromised data to exploit vulnerabilities in systems, for social engineering, or for performing reconnaissance into a target system. The advent of large scale analytics, AI, and IoT, provide new, much less well understood avenues for a malicious actor to have a cybersecurity impact through compromising the origination of data. We expand on this in our discussion, later in this document, around broader supply chain data risks associated with AI.

Cybersecurity and AI (and other) risks are greater than the sum of their parts

Cybersecurity does not exist by itself. We are pleased that this is explored in the CSF 2.0, which mentions that it should be included as an aspect of Enterprise Risk Management and leverage other risk aspects such as Financial and Privacy. However we note that many organizations still separate personnel with cybersecurity intuition and skill from those who are familiar with other business processes and even other IT processes. Cybersecurity, when considered alongside some other disciplines, presents unique combined risks that are greater than the sum of the parts and, thus, may not be visible to those who only consider one, or the other.

There is a very real risk that if each framework merely references the other, these “greater than the sum of their parts” risks will fall through the cracks, with practitioners of each discipline considering them another discipline’s problem. Instead, we feel that each framework should own these risks, while also pointing out that it is necessary to work with the other disciplines to properly manage them. In effect guiding practitioners to make an active decision in collaboration with other disciplines on which one is best placed to lead management of a risk where overlap exists, rather than passively assuming that it is handled by someone else.

Where the CSF 2.0 calls out other related disciplines, such as in Section 5, where the NIST AI Risk Management Framework (AI RMF) and Cybersecurity for IoT Program are referenced, we feel that to avoid them falling into the cracks, at least a high level identification and characterization of these “greater than the sum of their parts” risks belong in the CSF 2.0, at least in the Govern and Identify functions. Naturally, further elaboration of these risks and their effects on the other functions belong in a separate document.

In this section we will discuss some examples of these “greater than” factors in the context of AI, for the purpose of illustration, particularly as the AI RMF was mentioned in Section 5. We acknowledge that further detail probably belongs in a separate document. Cybersecurity

interacts similarly with other disciplines, such as IoT, and we feel should also be at least mentioned in a similar manner.

Defining AI

There are many different definitions of “AI” within both academia and broader society. In the general case, there is no one correct definition but rather, definitions that are useful in different contexts. In the context of discussing risk management, we find that there is benefit to using a narrower version of that defined in ISO/IEC 22989:2022 (and on which the definition in the AI RMF is also based). The definition that we use in the rest of this document is as follows, with our deviation from ISO/IEC 22989:2022 in italics.

“An AI (system) is an engineered system that generates outputs such as content, forecast, recommendations or decisions for a given set of human-defined objectives, *where the process used to generate the outputs cannot be practically (in the context of the application) derived or verified by humans using analytical methods.*”

We find that this definition provides a clean distinction from a risk management perspective, particularly because AI systems that satisfy this definition are much harder to risk-manage through (human) analysis, with the resources reasonably available. They are also more likely to behave in ways that may be useful, but may not be understandable. We feel that it also provides a more intuitive definition as it explains the “artificial” in artificial intelligence (because “human” intelligence could not practically create the process within the application context).

Note that our definition is sector and technology agnostic. In particular, it includes, but is not limited to, machine learning, statistical machine learning, neural networks, and deep learning.

AI presents special supply chain risks

It may not be immediately apparent to those performing cybersecurity analysis as part of the Identify function that the use of AI has the potential to invert the supply chain somewhat. Organizations more broadly may not be fully aware of the data that is flowing backwards through the supply chain, often without any human oversight. Even those who may be aware of the existence of such data, such as IT and cybersecurity practitioners, may not be fully aware of the need, and ways, in which this data should be risk managed from a cybersecurity perspective.

Examples of data that can flow in unexpected ways through and across the supply chain include:

- Telemetry from internal, 3rd party organization, and public facing apps and websites monitoring such factors as automated surveys, user engagement, and other analytics, and that are processed automatically into data that is consumed by business processes.

- Training and fault data that is automatically collected and analyzed from customer devices, such as apps, industrial machines, or self-driving cars. This includes such data as images, other raw sensor data, and core dumps.
- Chatbots that interact with the end users.
- AI systems that are used by business decision makers that may be influenced by external entities, such as translation systems, search engines, or websites (internal or external), particularly those that host external sponsored or advertising content.

We feel that the Govern function, Cybersecurity Supply Chain Risk Management (GV.SC) category, should more strongly acknowledge the potential for data, and cybersecurity risk, to flow through the supply chain in directions that may not be immediately intuitive or traditionally recognised. That data can flow down, up, and horizontally in the supply chain, and can occur without knowledge of the originator or custodian.

While these unusual flows have always existed, its increasing use in AI driven analytics and, in some cases, automated business process decision making, increases these risks and means that they may not be caught by traditional risk management techniques, such as human oversight. The key differentiator lies in our definition of AI, that the decision making is either opaque, or sufficiently complex, that problems may not be obvious or practically traceable, allowing malicious actors ample opportunity to hide their tracks. Furthermore, the sheer quantity and complexity of data can make it difficult or impossible for human judgment to be applied, or for transparent guardrails to be placed on the system. Where the AI systems can then provide data to both internal and external consumers, these can now surface as cybersecurity risks.

For example, a malicious actor may be able to manipulate or poison telemetry data that is flowing backwards through the supply chain in order to influence an AI system to make a particular business decision. Manipulating data to influence outputs of AI systems that are relied on for decision making has already been demonstrated. For example, a Berlin artist causing a ‘virtual’ traffic jam in Google Maps using a child’s trolley filled with cell-phones[4]. In this case, the manipulation was easily traced and admitted to. If a process is outsourced to a third party (say Software as a Service), or if AI is used, some of this transparency is lost but the risk for altering behavior on ‘bad’ data is still there. As another example, a generative AI system, such as a chatbot, may be used to exfiltrate data from an organization or 3rd parties, or may be poisoned as part of a (spear-)phishing campaign.

Perhaps this is also a role for the Identify function, Asset Management (ID.AM) category where this reverse flowing data may be generated by unsupervised assets that are within the organization, but that generate this data in response to 3rd party actors. Examples include IoT assets and virtual assets such as servers that interact with the public through web browsers, apps, and so-on.

As an aside, we do also feel that the language around the GV.SC category in general could be clearer and more consistent. In some places in the CSF 2.0 document, GV.SC appears to refer

[4] https://www.washingtonpost.com/technology/2020/02/04/google-maps-simon-weckert/

to just the supply chain relating to cybersecurity products and services (bringing to mind topics such as making sure that virus definitions are authenticated or that a 3rd party firewall vendor is appropriately trustworthy). In other places in the document it appears to refer to the analysis of cybersecurity risk across an organization's entire supply chain ecosystem (which we assume is the intended meaning).

AI presents special behavioral risks

We feel that one of the biggest risks that AI poses to an organization's cybersecurity posture actually has very little to do with the AI system itself, and everything to do with the nature of AI. An implication of our definition is that people are going to be more tolerant of the AI system doing things that they don't understand. After all, if the behavior was completely understandable, even to a lay person, it is unlikely that the "Artificial" in AI would be necessary. Furthermore, if the people in the organization have the impression, correctly or otherwise, that the AI systems used by the organization learns through time and can change their behavior, they may also become more tolerant of unexplained changes in system behavior.

This presents a significant and under-appreciated cybersecurity risk, for it means that one of the major resources for cybersecurity, people in the organization being vigilant for changes or unexpected behavior, becomes impaired. Worse yet, this risk extends not just to the AI system, but to other systems in the organization that may not have AI components at all, because their users may become used to expecting unusual or different behavior. This presents a much greater opportunity for malicious actors to hide their activities, including those that may generate some visible variation in system behavior. Inside malicious actors, in particular, could more readily use the excuse of "Oh it's just the AI system being strange".

This effect extends not just to the human aspect of the system, but also to automated anomaly detection. Learning what "normal" behavior means for a system, be it in terms of network traffic, decisions, or otherwise, becomes more difficult, particularly if the AI system changes its behavior through time, or in response to more complex inputs that the anomaly detection system may not be aware of.

We feel that an important first step to addressing this risk exists in the Govern function with the identification of desirable and 'must not' (guardrail) behaviors. It then flows to the Identify function, Asset Management (ID.AM) category, where systems with behavior that change through time, AI based and otherwise, without explicit, obvious outside control (e.g. not software patches), should be identified. It could be argued that this is part of ID.AM-08 regarding system, hardware, software, and service lifecycle, but we feel that this is a sufficiently unusual risk that it deserves its own subcategory. More so if aggregated with Third Party supplier risks if they are supplying, managing or operating the AI system, where there needs to be an element of trust and common interpretation to identify desirable and undesirable behaviors initially and dynamically in the life of the AI system.

There is also an opportunity to address the risks posed by AI on the behavior of people within the Protect function, Awareness and Training (PR.AT) Category. Here, it is important to not just perform general or specialized tasks with security risks in mind, but to also be aware of what may be considered “unusual” for different organizational systems. More generally, we feel that the PR.AT category may benefit from a subcategory relating to people being on the watch for unexpected behavior, even if it is not part of their role.

AI is special, along with everything else

Having spent the previous page talking about how AI risks are special, it is also important to remember that AI systems in particular, and “special” systems in general, are not special from the perspective of being immune to cybersecurity risk.

One concerning trend we have been seeing is for anything that includes “AI” in particular, and any “new shiny” technology in general, to be given special treatment by an organization, including circumventing governance, cybersecurity, and other risk management controls, sometimes on the instructions of non-technical management.

For example, one of us was recently in a presentation about the adoption of a particular new, important, technology in a safety critical field. The presenter, at one point, said to a mostly management audience “Remember to tell your IT folks to drop the firewall to make sure that the system can make its connection”. As another example, during a session discussing technology replacement, there was a prevalent preference for ‘cutting edge’ products over existing ‘tried and tested’ products. When the capabilities of the new technology were flagged alongside the additional risks that they would create, the feedback was to ‘ignore those’.

The CSF 2.0 of course provides exactly the kind of framework within which policies to address such risks can be developed. However, much of the commentary around its use, particularly in Section 3, reflects the use case of taking an organization, with its existing processes and procedures, and improving its security posture. There is some implicit recognition that new technologies, processes, and systems might exist insofar as the framework is depicted as cyclic, responding to changes in environment, policy, and technology. We feel that response to technology and new capabilities, as deliberate actions, should be called out more specifically.

It may be beneficial, within the discussion in Section 3.1 on Creating and Using Framework Profiles, to be more explicit in highlighting the use of the CSF 2.0 to evaluate not just the current or target cybersecurity postures of an organization, but also the expected posture of the organization in response to a non-cybersecurity change. This can include the introduction of a new technology, be it large or small, changes to the risk environment or changes to stakeholder expectations. Perhaps this even warrants the creation of a new profile type, perhaps called “Proposed New Technology and Other Changes Profile”, alongside and compared against the Current Profile and Target Profile. Vendors of new and, perhaps, less poorly understood, technologies can then also be encouraged to provide their own profile templates that outline the responsible integration of their system into an organization’s broader cybersecurity posture.

Cybersecurity has a human factor

Truly effective cybersecurity starts and ends with people. We feel that the commentary around the CSF 2.0 could benefit from being more explicit in its acknowledgement, and mitigation, of inevitable human failings and characteristics. We call out three specific examples in this section by way of illustration.

Avoiding shelfware

Organizational management is littered with examples of “check-the-box” mentality and “shelfware”, whereby an organization merely aims to satisfy a given standard that requires them to have a policy or a procedure, but that documentation sits on a shelf, never to be touched again. This is exacerbated by complex supply chains, with separate organizations, driven by different motivations and levels of risk tolerance, and where there may be an incentive to mislead with regard to risk posture. A common example is claiming compliance with a given standard, when in reality that compliance is limited to a specific, possibly inconsequential business process.

This problem is not limited to malicious intent. Even well meaning, but ill-informed, organizations can suffer from the creation of shelfware, such as by applying policies that are not suited to the organization or that are technically unactionable and, thus, ignored.

We feel that the CSF 2.0 can benefit from a greater acknowledgement that controls, policies, and procedures should be technically actionable and fully realizable within the business context. Monitoring and adjusting for this, we feel, is just as important as having the policy. An imperfect policy that can actually be implemented with the resources available may (although not always) be better than an ideal policy that cannot actually be implemented and is ignored by the people who are actually doing the work.

Perhaps this discussion could be incorporated into Section 3.1, regarding target profiles, to ensure that they are technically actionable and continuously updated to ensure that, when implemented, they actually do what is intended. Our previous suggestion on strengthening the link between the CSF 2.0 and other frameworks would help this, while also improving the visibility of the material practical benefit of the CSF 2.0 for business processes and activities.

Section 3.2 on assessment could also benefit from more explicitly calling on organizations to look for controls, at the implementation level, that are not performing as intended, be it because they are impossible to actually implement, or perhaps because there are unforeseen characteristics of the business process.

Section 3.4 on communications is a critical part of the CSF 2.0 that can prevent the generation of “shelfware”. We feel that its contribution could be better emphasized earlier on in the document and executive summary. In particular its mention of “bottom-up reporting” could be

expanded upon and emphasized, particularly in the context of ensuring that the CSF 2.0 doesn't just become a box ticking exercise.

We feel that Section 3.5 on managing supply chain risk could be more explicit about encouraging detailed communication between organizations on the supply chain, and the development of industry standard profiles that are sufficiently detailed so as to be technically informative, rather than becoming another "badge" that an organization can place on their marketing material, but is meaningless from a cybersecurity perspective.

Beware of phantom training

The aforementioned culture of box ticking can be even more problematic when applied at the personnel level. Much as the Protect function, Awareness and Training category, PR.AT-01 and PR.AT-02, talk about users being provided with training, we feel that it should include language such as "reasonable", "relevant", "actionable", or similar. It is all too common for organizations to deluge personnel with irrelevant training that they do not learn from or cannot apply, just so that a box can be ticked to say that they have had the training. A review process that applies metrics and records personnel cyber actions outside of training would be useful in gauging the effectiveness of training. The CSF 2.0 is an ideal opportunity to call this out, and to emphasize that the training must be relevant and technically actionable, rather than simply that an organization can tick this box by providing personnel with a potentially nebulous, irrelevant, generic training course.

Beware of undoing training

One human aspect of cybersecurity that we feel is missing from the CSF 2.0 is the interaction between cybersecurity training and other activities within the organization that may be counterproductive. The aforementioned PR.AT-01 and PR.AT-02 talk about user training to ensure that those in general and specialized roles have the "knowledge and skills" to perform tasks "with security risks in mind". This can easily be un-done if other business processes routinely require them to ignore these risks.

A common example concerns email phishing. Many organizations require employees, contractors, and associates to be trained to recognize and report phishing emails, with signs such as a spoofed "from" field, links to click on that do not go to organizational domains, look-and-feel that does not match the rest of the organization's branding, and language that seems designed to induce a sense of urgency.

Yet those same organizations will contract 3rd parties to, for instance, provide required training, auditing, or other services that mean that their employees will be receiving legitimate emails, from these legitimate 3rd parties, that may have spoofed "from" fields, links to click on, go to domains outside of the organization, have a different look-and-feel to the rest of the organization, and language that is designed to induce a sense of urgency.

It is little wonder, then, that phishing is still a common problem.

There does not really seem to be a natural place in the CSF 2.0 where this human aspect of cybersecurity truly belongs. We have suggested some possibilities below.

- In the Govern function, perhaps under the Supply Chain category (GV.SC), policies could be put in place to ensure that interactions between 3rd party suppliers and organizational personnel are consistent with risk management training. In the aforementioned email example, this could be a policy that resembles that of some banks, whereby their policy is to never send emails with links, but rather to request that the user log into the organization's secure site where they will find a way to authenticate and act on the message.
- In the Identify function, perhaps under the Risk Assessment (ID.RA) category, a human interaction subcategory could be added that encapsulates these risks, from both inside the organization and with 3rd party suppliers.
- Similarly, perhaps the Identify function could benefit from the addition of the concept of feedback from internal personnel as to perceived cybersecurity risks, such as receiving an actual legitimate email that looks like a phishing attempt. This might mirror ID.RA-08, which is more concerned with vulnerability disclosures (presumably from external sources).
- In the Protect function, perhaps as an addition to the aforementioned PR.AT subcategory, it may be beneficial to add ways in which personnel can be proactive about verifying information before acting on it. In the aforementioned email example, for instance, there could be a list of legitimate external domains and emails where legitimate emails could be expected.

There can be too much cybersecurity

There is the implication that the "current" profile has more cybersecurity risk compared to the "target" profile, and that "improvement" implies "reducing" cybersecurity risk. Even ignoring the fact that cybersecurity risk is multidimensional, and that it may be necessary to increase risks in some aspects while reducing it in others, we feel that the language and commentary around the CSF 2.0 may also benefit from acknowledging that it can be human nature to fear the unknown, and that sometimes this results in an organization that is too risk averse, particularly if there is little institutional cybersecurity expertise.

This is implied in Section 4.2 in terms of balancing cybersecurity risk with other risks, and we feel that this balance should be called out more prominently. This includes the acknowledgement that perhaps the target profile may in fact involve a greater level of cybersecurity risk than the current profile, if the organization was previously excessively conservative. This is very relevant for organizations outside Critical Infrastructure sectors where the level of acceptable risk can vary greatly and dynamically. In these organizations, application of cybersecurity that is not also dynamic can in fact hinder operations.

The CSF 2.0 provides an opportunity for such overly conservative organizations to evaluate their risk, and perhaps come to the conclusion that maybe their risk management can be loosened up and that they can say “Yes” to new capabilities. The aforementioned “Proposed New Technology Profile” may also be a way for an organization to convince itself that implementing a new technology will still mean that their organization’s cybersecurity posture is appropriate or, perhaps, even improve it overall.