17 March 2023

To: cyberframework@nist.gov

Subject: Comments regarding CSF 2.0 – Concept Paper, Virtual Workshop, and In-Person Workshops

Greetings, Overlords of the CSF:

Thank you for all of the great work that's gone into not only the CSF, but all the rest of the incredible pubs and products that NIST creates. The RMF, the FIPS library, the SPs and NISTIRs, the CSF, and everything else, is simply the greatest collection of information and cybersecurity goodness on the planet. You have my heartfelt thanks for all that you've done for us over the years and are still doing for us today.

I was fortunate enough to attend both the virtual and in-person workshops and have reviewed the CSF 2.0 Concept Paper. Overall, the proposed changes are relevant, timely, valuable, and they reflect the current cybersecurity landscape. Thanks for the opportunity to submit feedback, and here are some topics regarding which I would like to submit the following comments for your consideration.

1. **We should fully commit to the use of online and automated resources as much as possible for CSF 2.0.**

The rate of "interconnectedness" between the CSF and other references, and the volume of other references, is already substantial, and is going to do nothing but continue to grow. It is already intractable to sustain a list of informative references as a "publication" because of the turbulence of that dataset. There remains a gap between the outcome-based elements of the framework (theory level) and the application of controls extracted or derived from informative references (practical level).

A) NIST should proceed to showcase CSF 2.0 via the Cybersecurity and Privacy Reference Tool (CPRT), but NIST also needs to harmonize and disambiguate CPRT use with the Online Informative References Program (OLIR). While a security practitioner or academician can read and understand the content about CPRT and OLIR, thus discerning their intentions and use, the content can be a bit obtuse.

B) Right now, the CPRT and OLIR can be intimidating to the average CSF consumer. Yes, the interfaces are very functional, and they are generalized so that they can be used by anyone, but we probably need to develop some instructional guidance and use case examples for CSF consumers, and some kind of training wheels for folks who could potentially be overwhelmed by the power and depth of these tools. An integrated "CSF landing page" that hooks all of this together in a coherent way, with integrated help and tutorial linkage would be most welcomed. We want the automation to accelerate and simplify the adoption and use of CSF, not boggle the minds of the poor souls who are trying to find the crosswalk between their industry standard and the CSF.

C) CSF 2.0 should remove every informative reference and use only the OLIR. The CSF 2.0 is going to come out as a document, and informative references should not be in it. Those references can go out of date within minutes of the publication of the framework and there's no point in expending effort on the "what's in and what's out" editorial decisions. Just don't publish the references. Instead, point everything to OLIR, now and forever.

D) While CPRT and OLIR are great starting points for improving the use of automation in managing the CSF (and related) content, NIST should further extend this across the Security Content Automation Protocol (SCAP) specification suite and the Open Security Controls Assessment Language (OSCAL). There has been significant investment in SCAP and OSCAL, and they remain underutilized. Wherever possible, NIST should cite SCAP/OSCAL capabilities, and provide SCAP/OSCAL content, to address elements of the CSF. One of the jewels in the CSF crown is that it's vendor and technology agnostic. SCAP and OSCAL are the same. There are a lot of

correlations between CSF outcomes and the expressive abilities of the SCAP specifications, from Asset Reporting to Checklist Configurations. While download of CSF materials in JSON, CSV, and XLS formats are supported, so should download of materials in XML and other SCAP formats. While NIST is pointing out the relationships between CSF 2.0, CPRT, and OLIR (NISTIR 8278), any mention of the SCAP NISTIRs and SPs is notably missing. This should change and CSF should help further the adoption of SCAP because of the neutrality, agnosticism, and open-source nature of the specifications.

E) OSCAL continues to be a domain of Dark Magic, at least as perceived by many of the uninitiated. Actually, it's pretty much Dark Magic no matter who you're talking to, but I digress. The layers, models, and Metaschema are well-designed, elegant, and powerful. But they are intimidating and most of the NIST documentation is oriented to the more technical readers. NIST should place greater emphasis on the development of content flavors like "OSCAL for the CIO," "OSCAL for the IT Manager," and even something like "OSCAL for the Small Business Owner." Most importantly, we need "Use of OSCAL in Support of the CSF." Right now, the use cases for OSCAL cite "Managing Multiple Regulatory Frameworks," and that should be resonating strongly with CSF and the OLIR content, but there is not even a mention of CSF in the OSCAL collateral materials. There's strong emphasis in the OSCAL content on the NIST RMF (for quite logical, historical reasons), but we need equal or greater emphasis on NIST CSF. There's an OSCAL System Security Plan model, but we need OSCAL CSF Profile and Tiering models. The NIST and the FedRAMP PMO are maintaining OSCAL content on GitHub, but there's no content for CSF users. This needs to change. Not only should the CSF Core be encoded for OSCAL use, but NIST should strongly support an OSCAL/CSF Community of Interest (COI) where Adopter's Workshops, OSCAL 101 Workshops, training, CSF OSCAL promotional materials, etc., are produced and sustained. **SCAP and OSCAL remain powerful but underutilized tools that, in concert with the CPRT and OLIR tools, could significantly improve the CSF ecosystem and quality of life for CSF consumers. They need to be brought to the forefront of CSF use, not left on the sidelines.**

F) NIST should promote the development and use of open-source tools that help with the use of the CSF. It's understood that NIST cannot endorse or recommend products, and that NIST must remain a neutral, disinterested party in such matters. But NIST is already publishing a "Resources" page with pointers to third-party content (https://www.nist.gov/cyberframework/framework-resources), and there is an open solicitation for the submission of "risk management resource" content at https://www.nist.gov/cyberframework/resources/risk-management-resources . This should be significantly enhanced, and a more robust, community-centric model should be adopted. Via other online channels, we see resources like the CSF Tools and OpenSCAP sites, but they are not found as CSF resources, nor are they likely to ask to list themselves there on the NIST page. The community should be able to point to these sorts of resources and discuss them. NIST should encourage their development.

G) Right now, the submission method for adding content or a link to the Framework resources page is to send a description or content to the cyberframework@nist.gov e-mail address. There's no disclosure of the approval process, nor is the community asked to participate in the process. There is no transparency here. While NIST is the Overlord of the CSF, it's the community that's going to create and manage ways to implement it. Instead of using an anonymous e-mail address and opaque approval process, NIST should allow the community to upload content, discuss it, up-vote/down-vote it, and otherwise self-moderate it.

Before everyone freaks-out and assumes that this will result in a *CSF Reddit or Twitter Hellscape*®, we should look at the model being used in the open-source community. Open-source software (let's just roll with the Linux project since it's the most famous one) can self-police, protect itself from evil, and continue to develop and flourish through a community-based system. There is an Overlord for the Linux kernel (Linus Torvalds), but there's also a community of developers, testers,

approvers, and participants who are simply trying to contribute and *Do Good Things*. One need only look at Github to see how a community-based effort can yield results of incredible value.

The same can be fostered around the CSF. While NIST will remain the Overlord and wield the power to smack-down bad actors, most of the work of developing, reviewing, approving, and maintaining CSF resources should lie with the community. NIST cannot fund the horsepower needed to truly grow and sustain an ecosystem of CSF resources on its own. Let the community help with that. There are numerous successful models of self-managing, open-source communities that can produce great things, and we should at least try to build some for the CSF.

H) While we don't need to start reviewing CSF profiles on Yelp, some kind of review and rating system should be put into place, with the ability to provide feedback and discussion on CSF-related content.

I) CSF Profiles are a valuable facet of the CSF, but they remain an abstraction until they can be measured or assessed against. Like OLIRs, CSF Profiles should have a "lifecycle," where they progress from Work-in-Progress, Draft, and then Final status. CSF Profiles that utilize any references to other issuances and are submitted only in the form of written prose should remain in the Work-in-Progress and Draft stages. In order for a CSF profile that references any other issuances to reach "Final" status, it should include a working and vetted set of OLIR content, including the Rationale and Relationship settings. Right now, the majority of the CSF Profiles (https://www.nist.gov/cyberframework/examples-framework-profiles) contain only a basic linkage between the CSF subcategories and the informative reference, and while some do express the extent of the mapping, the terminology definitions and use are not consistent. Everyone needs to get onboard with using not only the OLIR system itself, but the terminology and definitions found in NISTIR 8278A. NIST should also place an emphasis on, and preference for, profiles that are built with SCAP and OSCAL content.

J) Although not CSF-specific, the proposed changes in NISTIR 8278A Rev 1 are most welcomed, particularly greater clarity and specificity in the IR lifecycle. While the new content more thoroughly covers the "end of life" for an OLIR, adding a "Deprecated" status to the lifecycle would be beneficial.

2. Yes, there should be a Govern Function

The Govern function should span the entirety of the CSF lifecycle, like the PM control family in NIST SP 800-53. The CF Concept Paper, Section 4.1, seems like a solid approach to adding this Function, but there are concerns that this function could be intimidating or confusing if not approached correctly.

If we look at the NIST Privacy Framework, for example, we see that the basic Govern function is "*Develop and implement the organizational governance structure to enable an ongoing understanding of the organization's risk management priorities that are informed by privacy risk.*" Yes, that's a fine function definition, but it's not one that's going to properly inform Small-to-Medium Business (SMB) without clarification. As soon as we see "organizational governance structure," we've got a lot of entities that are going to balk or they're going to get confused. We need to simplify the notion of Governance to the most basic of definitions:

> "Governance is how the organization determines priorities and then allocates resources to those priorities. It's the process by which the organization is controlled and operates."

It is critical that the CSF explains that Governance is already happening in every organization, regardless of size, from a multi-national corporation to a one-person company operating in a garage. Somebody, or some *thing*, is deciding what the organization is going to do, and then it's allocating resources to do it. That's Governance. You do not need COSO and COBIT, and you don't need a "governance structure" or some "governance department" to effectively implement good CSF

Governance, although there are some organizations that use those things. If we skew the CSF more towards the larger entities, we're going to alienate the SMBs, startups, and others who are trying to implement CSF as they grow. It is important that the Governance Function **is an enabler of CSF implementation, not a "barrier to entry"** for organizations that don't have formalized governance structures. You can do good governance without a big, complicated "governance machine" in your organization.

The bigger entities that are already doing more formal Corporate Governance already know that they're using COSO and similar things to operate. They don't need the CSF to tell them that they're looking at instantiating the Functions into their "organizational governance structure." The CSF needs to tell the smaller and less-mature organizations, "Hey, you can still do this – it sounds complicated, but it's really just a matter of scale." It's the SMBs who need more handholding, not the big corporations or agencies.

3. Vendor and Technology Neutrality should remain a key principle of the CSF, and we should add neutrality regarding Organizational Size and Structure

Everything in Section 2.6 of the Concept Paper is Really Good Stuff®. The proposal that "Additional guidance on tailoring for specific technologies or applications may be best accomplished by CSF sample Profiles, mappings to specific standards or guidance, or implementation examples," is a great position, and adding Playbooks to that list would be welcomed, whether they're called "Playbooks" (like in the case of the AI Framework) or SP 1800-series publications, or NISTIRs.

Statements regarding organizational size neutrality occur only twice in the CSF: Once in the Executive Summary (*The Framework enables organizations – regardless of size, degree of cybersecurity risk, or cybersecurity sophistication – to apply the principles and best practices of risk management to improving security and resilience*), and once in Section 1, the Framework Introduction (*It is intended to be useful to companies, government agencies, and not-for-profit organizations regardless of their focus or size.*). The rest of the CSF remains mute on the topic. There is no specific mention anywhere in the CSF of scaling the framework or the scalability of the framework.

CSF 2.0 should improve upon this. The CSF should continue to declare that vendor and technology neutrality are cornerstones of the Framework, and organizational size and structure should explicitly be added to that. A CSF section that specifically addresses scaling the Framework, and its adaptability regardless of organizational size and structure should be clearly present.

A) Section 3.2 of the CSF does contain content about how an organization can use the Framework, and it does cover steps to "Prioritize and Scope" the effort, and then to "Orient" the effort. While these could be construed as steps that somewhat bring about scaling the Framework, they are too vague in this regard. Scoping and scaling are not the same thing, for example.

   While there was some (mostly breakroom) discussion during the workshop about the steps in CSF Section 3.2 being moved to the new Govern function, they can't simply be transplanted there. However, the general notions in 3.2 certainly can inform the new Govern function.

   The new Govern function should have a "Scale the Framework" subcategory or some other scaling component in the Function but needs to remain agnostic as to the process of scaling.

B) There are similarities between Scrum and the CSF in that they are both frameworks, and both are meant to be simple, elegant, and durable. They are outcome-based (and/or "value-driven" depending on how you want to spin the nomenclature). Scrum, the leading example of an "agile" framework, clearly showed that it could increase value in organizations and to customers, has been considered "limited" because of the small scale inherent to the Scrum philosophy. Large organizations seeking to adopt Scrum have been challenged in how to scale it to address large projects, large teams, multi-team projects, and how to function with enterprise-level governance,

compliance, and yes, security. Significant work has been done on how to address the scaling issue – Projects such as Scaled Agile Framework (SAFe), Scrum@Scale (SaS), Large Scale Scrum (LeSS), and Disciplined Agile (DA), are only a few of the more popular approaches.

A future in which various approaches for scaling the CSF can be envisioned. As in the case of Scrum, the source itself should remain small, simple, elegant, and durable.

4. This won't be popular, but we should dial-back on the C-SCRM in the CSF

We're on a slippery slope with the Cyber Supply Chain Risk Management (C-SCRM) content. CSF 1.1 responded to very valid needs and requirements regarding supply chain risks. But it also cracked open the door to potentially inviting too much C-SCRM detail into the Core. We are all aware of the fact that there is a large amount of feedback and requests to add even more C-SCRM content, but we're running the risk that the CSF will become too "C-SCRM heavy" and that C-SCRM will begin to overshadow the other categories and sub-categories.

More importantly, we're looking at a situation like that of the Governance Function: If we add more complexity and detail, particularly the kind of detail that larger organizations seek, we run the risk of alienating smaller organizations that could really use some basic C-SCRM guidance but are overwhelmed by all the jargon and detail that is seeping into the CSF.

We need to keep the CSF Core small, simple, elegant, durable, and agnostic as to the size of the organization implementing it. If we skew the CSF Core towards the C-SCRM needs of the big entities, we're doing a disservice to everyone else.

Yes, there needs to be more C-SCRM content and, yes, the big entities need more specific guidance on large-scale C-SCRM activities, but the CSF Core is not the place to do that. The Core should have only the most basic of C-SCRM outcomes, and the more detailed and complex C-SCRM content should be provided elsewhere. It's noted that the Concept Paper states that respondents don't like the idea of developing a separate framework for C-SCRM, but what about a CSF C-SCRM Playbook? The AI Risk Management Framework has a draft AI Risk Management Framework Playbook, and that sems to be a model for what can be done with the CSF. Just as the AI Framework Playbook "suggests ways to navigate and use" the AI Framework, the CSF Playbook could do the same. If not in a CSF Playbook, perhaps a separate "CSF C-SCRM Playbook: Scaling Cyber Supply Chain Risk Management for any Organization?" An SP 1800 series implementation guide, or a NISTIR would be other good options.

The bottom line is that C-SCRM can be really big, complicated, and intimidating, and that's the exact opposite of what we need in the CSF Core. But we need it somewhere.

5. Please don't put any examples, notional or otherwise, into the CSF itself

This is another situation where there's a solid and justified demand for CSF-related content, but it's best not delivered via the CSF itself. As Spock so astutely taught us in Star Trek, "*Having is not so pleasing a thing, after all, as wanting. It is not logical but often true.*" The CSF already offers us so much, but we want so much more.

There is significant value to be had in adding examples, but doing so is also fraught with risk. For every example NIST includes, there's five other examples that are just as good that won't be included. Every example will also unintentionally limit the thinking of the reader, no matter how much such limitations are warned against, that the example is "how" to "do" the subcategory. We can all envision the meetings where some forward-thinking practitioner devises a perfectly useful means of reaching a CSF outcome only to be thwarted by the reply, "But that's not how NIST does it in the CSF example." Even with the best of intentions to prevent it, the examples will be seized upon as a baseline by organizations that simply don't know what else to do. Organizations will use the examples as their implementation model, only to be disappointed when the example they use is later removed from the CSF because maybe it wasn't the best example and it gets revised, and now they're going to go through the

machinations of changing their implementation *simply because an example was changed*. We all know it's going to happen.

Put the examples into a Playbook or, better yet, devise something akin to the OLIR where various examples can be crafted and shared. Let the community contribute, and self-referee. Let the examples grow, evolve, and retire as CSF implementation continues to spread and improve. Niche use cases should have a home where their "best of breed" examples can be published and shared, and not limited by pre-packaged examples in the Framework itself.

6. We need to carefully address the Cybersecurity Measurement and Assessment Needs

It's entirely understandable that the community wants guidance about how to measure and assess use of the CSF. Depending on the guidance that is given, however, this could become a muddy and difficult domain to understand and manage.

A) The Implementation Tiers remain a domain of Dark Magic, and further guidance on how to use them in measurement is definitely needed (actually, just using them in general). The fact that they are not maturity model levels, yet look remarkably like maturity levels, combined with the fact that many organizations are using maturity models as part of their organizational risk management and operational capability progression and measurement, makes for a unique, complex, and potentially treacherous landscape. The CSF should include guidance about how to use the tiers alone, and in concert with maturity models, so as to extract optimal return on investment. Organizations (and their assessors and auditors) can easily become overwhelmed with so many measurements and assessments, with overlapping drivers and constrainers, that more resources are spent *measuring* things than actually *doing* things. How to simultaneously blend all of this so that it's actually productive is a domain where further investment should be made.

B) The NIST National Cybersecurity Center of Excellence (NCCoE) and the U.S. Department of Energy (DOE) Office of Cybersecurity, Energy Security, and Emergency Response (CESER) deserve a hearty "Thank you" for the great work that they've done on the Cybersecurity Capability Maturity Model (C2M2) and the mapping to the CSF. With the recent March 9, 2023, announcement of the updated draft mappings, it's great to see the refined forward and backward mappings to and from the two focal documents, as well as the C2M2 to CSF Tier mappings. Further, the fact that well-formed OLIR content has been published shows that mapping efforts using the standardized NISTIR 8278 expressions can clearly and usefully can be employed in the mapping process. This effort should be showcased as an "Example of Success" - This, folks, is how you do it.

C) We need to differentiate (carefully and very deliberately) between assessing the use of the CSF itself and assessing the "controls" that are employed by organizations in support of the CSF outcomes. While the CSF provides a "taxonomy and lexicon to communicate the outcome of [an organization's] measurement and assessment efforts," it very correctly does not specify methods and actions that bring about the CSF outcomes. It's important that methods and actions to measure and assess (the methods and actions used to bring about outcomes, aka, the "controls") are also not specified, and that CSF users are clearly informed that any measure or assessment of CSF implementation is not the same thing as, nor is it in any way a replacement for, the appropriate measurement and assessment of the "controls" that are used to deliver outcomes. This gets all very meta because at some point the use of the CSF itself can become a control, but the CSF needs to remain at a level of abstraction that is not prescriptive about the controls that deliver the outcomes. The statement in the Concept Paper that "NIST will not put forward a single approach to assessment in the CSF 2.0 in order to continue flexibility in how organizations may implement the Framework," is spot-on and a critical element of any effort to talk about measurement and assessment.

D) The intent to include examples about how organizations have used the CSF is lauded, but is this better delivered in supplementary materials, such as a Playbook or an SP 1800 issuance? As we seek to **keep the CSF simple, elegant, durable, and agnostic as to the size of the organization implementing it,** it would seem that this rather hefty topic deserves more than just a section in the CSF document. To fully address it in the CSF could overload the CSF publication with a large amount of content, and to over-simplify the topic into a summary will not do it justice.

E) Can the OLIR be expanded to allow for assessment and measurement methods to be linked at the same time the control requirement is mapped? For example, the OLIR mappings between the CSF and SP 800-53 already exist (in both directions), but the Rationale, Relationship, Fulfillment, and Strength of Relationship content are noticeable missing (which I hope can be remedied, please). We already know that there's a strong binding between the controls in SP 800-53 and the assessment methods in SP 800-53A. Does the same relationship exist between a CSF subcategory and the SP 800-53A assessment method, if the CSF subcategory is mapped to the corresponding control in the SP 800-53? If not, why not? If we map a subcategory to an SP 800-53 control, what confidence do we have the SP 800-53A assessment method is also applicable? If the assessment method won't work, is the mapping to the control valid in the first place? The 800-53 and -53a is just an example, but these kinds of correlations exist elsewhere, and we need to address how this can work.

   It would seem that we should be able to map assessment methods in the OLIR just the same as we can map other informative references. When we say that control XYZ out of some informative reference is "EQUAL" and "FULFILLED," with a STRENGTH of 10, to a subcategory in the CSF, and there's a known assessment method for control XYZ, then what do we know about the equality, fulfillment, and strength of that assessment method in relationship to the subcategory of the CSF? If the answer is, "not a damn thing," then it would seem we need a way to map assessment methods totally independently of the informative reference itself (but still use OLIR to do so). If we can start with an assumption of "there's probably some correlation, but don't assume the assessment maps as strongly (or as loosely) as the control itself," then those might be "modifiers" applied to assessment methods that hook to controls that hook to subcategories, all via the OLIR. Oh, and don't forget this should all work with OSCAL and SCAP so we can automate this as much as possible.

   I'm sure, right now, that the Overlords and Data Wizards who do the whole OSCAL, SCAP, and OLIR projects, data structures, taxonomies, and schemas are reading this and all going, "Can we please just shoot this guy?", but efforts in this area could eventually reap significant rewards.

F) While the DoD-specific Cybersecurity Maturity Model Certification (CMMC) is focused on Federal Contractor Information (FCI) and Controlled Unclassified Information (CUI), efforts to harmonize and reuse elements relevant to the CSF should be promoted and undertaken. While it's clearly understood that NIST doesn't direct the DoD in such matters, the fact that they're working so specifically with the SP 800-171 Rev 2 and SP 800-172 issuances seems to express an opportunity for further well-crafted mapping and OLIR content to be published. Since the OLIR for the CSF 1.1 focal document to SP 800-171 Rev 2 is published, the reverse mapping should be attainable and that should further be able to support some CMMC mappings. Perhaps the NCCoE and DOE folks could host a happy hour for the DoD folks and help them understand how to do this correctly. Then everyone can reuse their CMMC content, and we can further both CSF and SP 800-171 Rev 2 Goodness for everyone involved.

7. International Collaboration is great, but can we keep in mind the mission and goals of the Department of Commerce, please?

It's understood that many organizations would benefit from the international use of the CSF, and that it would improve the efficiency and effectiveness of their cybersecurity efforts. Many of us who work in

cybersecurity are scientists, engineers, and academicians. The furthering of human knowledge, the sharing of research and intellectual advancements, and the continued advancement of human civilization are all noble goals that we share. NIST is to be applauded for the forward-thinking objective of prioritizing exchanges with foreign governments and industry as part of the CSF 2.0 development and furthering the recognition of the CSF as an international resource.

But NIST does not have infinite resources. Priorities must be decided. NIST is part of the Department of Commerce, and the Department's mission is "to create the conditions for economic growth and opportunity for all communities." The Department's overarching goal is to "Improve America's Economic Competitiveness," with one strategic goal being to "Drive U.S. Innovation and Global Competitiveness." The Department's Strategic Plan talks about goals to "Accelerate American Leadership," and "Strengthen U.S. Economic and National Security," including "Enhance the Nation's Cybersecurity." Of course, it's understood that by improving international cybersecurity, America's strategic goals are also supported. But with resources being limited, and with everything else being equal, it's important to keep in the mind that there's still work to do that's important to the U.S. The CSF started in response to a need to improve the cybersecurity of our Nation's Critical Infrastructure. It's great that it's grown far beyond that, and it's proven useful in so many other ways, but we should not lose sight of why all of this started in the first place.