MP170274

MITRE PRODUCT

MITRE

Best Practices for Technical Standard Creation

Dept. No.: J83D
Project No.: 25AOH830-JM

The views, opinions and/or findings contained in this report are those of The MITRE Corporation and should not be construed as an official government position, policy, or decision, unless designated by other documentation.

Guidelines for the Design, Socialization, Formalization, and Adoption of New Technical Standards

Approved for Public Release; Distribution Unlimited. Case Number 17-1332

Author: Charles M. Schmidt

Bedford, MA

April 2017

Abstract

One important enabler of a comprehensive cyber defense approach is uniformity of practice. Uniformity of practice reduces complexity and ambiguity in multi-component environments. It can also enable concentration of cyber defenses in areas where such practices can be most effective. Technical standards support uniformity of practice through a precise description of that activity. There is an ongoing need to develop new technical standards to reflect the continuous creation of new practices and technologies.

However, the development of technical standards is fraught with challenges, many of which can be unintuitive to those new to the process. This paper identifies major factors that can impact a standardization effort's chances of success and offers some guidance to address such roadblocks. It includes guidelines covering activities before, during, and after a standard is developed, as well as advice on navigating the cultures and operations of standards development bodies. Special emphasis is given to the role the US government can play in these activities. This guidance is based on input from over a dozen experts who, combined, have over a hundred person-years of experience working in standards bodies. Together, this information can improve the chances of successful engagement in a standardization effort.

Key Words

1. Cybersecurity
2. Technical Standards
3. International Standards Organizations
4. Industry Consortia
5. US Government

This page intentionally left blank.

Executive Summary

This paper looks at the important role that technical standards can play in securing enterprises, with a special focus on how and why the US government can play a role in the creation of these standards. Technical standards support the unification of practice with regard to technical activity through a precise description of that activity. For example, a standardized representation of some class of information provides a common means for different applications to express and potentially exchange that information. Technical standards offer many advantages to defenders. For example, if applications and services all utilize a particular network protocol or data format, then defensive tools can focus on monitoring that network traffic or data for signs of malicious behavior. By contrast, if applications employ custom protocols and formats, then each requires a separate security monitoring and enforcement capability. That makes defense more complicated (number of solutions required) and more expensive (cumulative costs) for administrators. Security standards (a special type of technical standard focused on cybersecurity activities) can allow security tools to share information and coordinate activities, possibly automatically, to detect and respond to threats. In these and other ways, the use of standards in enterprise applications and services makes the enterprise easier to manage, monitor, and defend. There is a constant need for new technical standards for IT systems. Whenever a new technology or trend leads to a flurry of new products, there are vendors who create their own, proprietary techniques to address critical questions of implementation. For the reasons described above, these proprietary techniques can complicate the task of monitoring, managing, and defending the enterprises that include these devices. The Internet of Things (IoT) and Cloud Computing, to name just a few recent IT trends, both provide plentiful examples of this behavior, with many incompatible products currently on the market.

Many of the technical standards used today evolved over decades of experimentation and revision. Given the rate of innovation in IT, it is unsurprising that, for many newer technologies and capabilities, standards are in short supply and that many products today rely on ad-hoc or proprietary protocols, formats, and interfaces. While existing IT standards can be, and generally are, employed in new products and branches of IT, many new products and trends bring with them unique operational constraints and use cases, which might not be addressed by current standards. For this reasons, many product vendors are forced to develop custom solutions in order to get their products to market. While these custom solutions put new technologies into the hands of consumers quickly, the resulting lack of uniformity across similar items in a class of products, or across classes of products in a larger category, complicates security for enterprises that use such devices. While new standards will not immediately remove this complexity (due to presence of legacy devices), as new products adopt more standards-based operations this will gradually reduce the complexity and improve the automation and security capabilities associated with these products.

The creation of technical standards is a challenging process and contains numerous pitfalls that can hobble those unaware of how best to proceed. At their simplest level, technical standards attempt to encapsulate, in a relatively concise document, the details necessary to enable uniformity of practice across a diverse range of implementations. While standards creation shares many aspects with other technical engineering efforts, its objectives impose unique requirements and challenges. This paper contains input compiled from available research as well

as conversations with over a dozen veteran practitioners in the development of technical standards. The result is a set of recommendations to navigate some of the greatest risks to successful standardization efforts. These recommendations are summarized as follows:

- A participant in a standardization effort should have a clear understanding of what constitutes a successful outcome of their engagement in that effort, while recognizing that other participants might have different, but not necessarily incompatible, definitions of success.
- When possible, convene a workgroup within a standards body around a draft solution rather than around a problem.
- Avoid having an initial solution "rubberstamped" by a standards body. Instead, use engagement with a standards body to evolve an initial solution so that it better meets the needs of a broad set of stakeholders.
- In any standard development effort, attempt to recruit engaged participants representing a diverse set of end users, vendors/implementers, technical solution experts, and standards experts.
- Agree upon and ensure that sound engineering practices are observed in the development of technical standards.
- Understand the values of the standard's target market and align design decisions and tradeoffs with those values.
- Avoid attempting to codify a high degree of detail in support of a highly diverse community. Detail in a standard should be balanced against the diversity in practice of the communities it serves.
- Be conservative in the number of use cases and features (optional or required) a standard supports.
- Understand how a standard supports customized use by vendors and users. If appropriate, support customization via well-defined extension points and procedures.
- Smaller, more narrowly scoped standards have a greater chance of successful completion and adoption than large, multi-part standards. If possible, start with the former and, if desired, build the latter through the composition of those smaller component standards rather than attempting to engineer directly to the larger vision.
- Be prepared to support efforts to raise awareness, advocate for adoption, facilitate interoperability, develop reference implementations and other resources, and otherwise provide support of a new standard for a period of time after the standard is published.
- Before proposing a new standardization effort in a standards body or consortia, evaluate the group to verify, among other things, that the group is a good fit for the proposed work, that the group can bring appropriate participants to the process, that the publication of the proposed standard by the group will impact the desired audiences, and that the procedures and culture of the group are acceptable.

- For standards bodies in which participant reputation and social credit play a role, actively seek to build this credit and consider supporting long-term participation in that body as a way to improve the participant's effectiveness in supporting efforts of importance.
- Actively build support for proposals within standards groups rather than relying on the technical merit of that proposal to speak for itself.
- Organizations seeking to represent their interests within a standards body should select participants who have a mix of technical competence as well as diplomatic skills to improve the reception to proposals from that participant.
- Allocate sufficient resources to actively participate in the writing, research, editing, and/or prototyping work undertaken in the course of a standardization effort.
- Expect that any standard creation effort will take longer than scheduled and allocate an appropriate level of resources.
- When actively participating in the development of a particular standard, allocate sufficient resources to allow physical attendance at in-person meetings of the standards body.

Standards development invariably involves multiple parties and, as such, always includes factors beyond the control of any single participant. For this reason, no list of guidelines can guarantee success, nor will deviation from one of these guidelines necessarily doom an effort to failure. However, following the recommendations provided in this paper helps participants reduce risks commonly encountered in standardization efforts.

Acknowledgements

The author would like to thank the numerous members of standards development teams, both formal and informal, with whom he has collaborated over the past several years. The lessons learned from those experiences allowed this document to be written.

Special recognition is owed to the following people who provided recommendations and input specifically for this paper: From MITRE, Jon Baker, Sean Barnum, Dr. Suresh Damodaran, Matt Hansbury, Dr. David Mann, Bob Martin, Bede McCall, Peter Sheingold, Kevin Smith, Stacey Stanchfield, and John Wunder. From outside MITRE, Lisa Lorenzin, Ira McDonald (High North) and Steve Klos (1E and TagVault.org). Their many suggestions and perspectives were greatly appreciated. Any errors remaining in this document are solely the responsibility of the author.

Finally, the author would like to express his appreciation to the US Department of Homeland Security's Front Office for the project that produced the original report on which this paper is based. Thanks also go to the MITRE Corporation's Cybersecurity Technical Center, for providing the additional funds needed to turn that report into this document as it currently exists.

Table of Contents

This page intentionally left blank.

1 Context

Technical standards support the unification of practice with regard to some technical activity through a precise description of that activity. This paper argues that the creation and support of technical standards are necessary to make IT security manageable. Technical standards can involve communication protocols, data formats, processes, architectures, and similar aspects of cyber systems. Without standards, this paper contends that the resulting use of custom and proprietary behaviors in IT products would require an operationally and economically unsustainable number of customized security solutions.

The US government has a compelling interest in advancing the state of IT security through technical standards. It also has a great deal that it can contribute to standardization efforts due to its experience and subject matter knowledge. The role of government provides it with perspectives and insights that commercial enterprises do not have. To this point, a recent AFCEA report notes, "Government plays a unique, but incredibly valuable role in facilitating – but not owning – this portion [standards development] of the mission." (Folk, et al. 2015) For these reasons, this paper argues that government support and involvement in IT standardization is an important part of its larger mission to enhance the cybersecurity of the US government operations, critical infrastructure, and the nation as a whole.

However, the creation of successful technical standards is challenging. In the past, many attempts to create standards failed or produced a standard that had little impact on commercial products. A range of factors influence how (or if) a standard is developed and adopted, including market forces, vendor buy-in, and individual personalities of participants, as well as technical aspects of a proposed standard itself. When dealing with new technologies and trends, the novelty and diversity of uses to which this technology is applied further complicates the process of creating standards that meet the practical need of implementers and adopters. To assist readers with navigating these challenges, this paper includes a discussion of risk factors and best practices for engagement with standards bodies. It also includes several recommendations that can help improve the chances of success of a standardization effort.

Audience

While the basic concept of a technical standard is quite simple (i.e., a technical standard supports the unification of practice with regard to some technical activity through a precise description of that activity), there is a great deal of complexity in bringing this concept to reality. Some of this complexity might not be readily apparent to those without prior experience in standards. One of the goals of this paper is to provide a set of recommendations that can help those new to standards development better prepare for and operate more effectively within a standards development effort.

The US government is not new to the topic of standards development. There are many examples of successful US government involvement in standardization efforts, and many government employees who have been valuable participants therein. These people are hardly new to standards development, but they too should derive some benefit from the recommendations presented in this paper. At minimum, it offers a source of consolidated information and references that they can share with those working in their areas of interest who are less experienced with standards. The paper may also offer practices that these practitioners might not have considered. Moreover, since this paper provides reasons for the recommendations it contains, it strives to provoke new ways to think about already adopted practices. Finally, this paper can help experienced practitioners defend good practices to managers or colleagues with less experience, who might see those practices as unnecessary or inefficient.

The recommendations in this paper cover activities prior to the convening of a standardization effort, the process of standards creation itself, and activities after a standard is published. The recommendations are intended to help the reader avoid pitfalls and situations that have been observed to increase the risk that a standardization effort will fail. By definition, any technical standard involves multiple parties (a “standard” used by only one party is really just a local customization), which means that no one group is able to single-handedly eliminate all chances of failure. However, by observing the recommendations presented in this paper, a participant can increase the value of their contributions to the group and help reduce some of the risks the effort faces.

Sources

In addition to looking at existing research on standards and standardization efforts, the author held discussions with over a dozen subject matter experts, from within The MITRE Corporation and from commercial industry, who have extensive experience in both the development of technical standards and participation in standards bodies and industry consortia. MITRE personnel have engaged in dozens of standards development efforts, usually under government sponsorship, and the MITRE subject matter experts who provided input to this paper have direct experience in the intersection between US government interests and standardization efforts. Together, these experts represent over a hundred person-years of experience in standards development across over a dozen different organizations. The practices of standards development vary across different standards groups, as well as between individual participants in such groups, and each of these experts has their own approach to standards development. However, despite these variations, there was broad agreement about many factors and practices and how they influence the success of a standardization effort. The recommendations presented in this paper reflect the general consensus of these subject matter experts.

2 Technical Standards and the Securing of IT

Protecting any connected computing device against remote attack is a significant challenge, as evidenced by the regular news reports of data breaches in corporate and government networks. Modern enterprises are distributed, complicated, and highly connected, all of which complicates cyber defense. Standardization of practice is one way of reducing some of this complexity.

Uniformity of practice allows defenders to concentrate their efforts on a small number of defensive products and tools that then are able to protect multiple types of devices, applications, and functional roles within the enterprise. Without uniformity of practice, it becomes necessary to support a separate solution for each type of enterprise component. For example, consider how difficult network traffic monitoring would be if every application used its own network protocol. One way to enable uniformity of practice is through the use of technical standards. By reducing the diversity of protocols, data formats, and practices used within an enterprise, complexity is reduced, attack vectors are decreased, and security investment can be concentrated where it will create the greatest return.

In addition, security standards, which are standards that focus specifically on security practices, make it easier for security products from different vendors, as well as custom-built tools, to interact with each other as part of a more holistic security solution. Without such standards, each defensive tool effectively becomes an information and capability island, limited in what it can observe and what it can share. This not only means that such security tools are unable to utilize wider context in their understanding of the information and events they collect, but that it is up to the security analysts to manually bridge these information islands, necessarily a tedious and time-consuming task, prone to human error. By contrast, if these same tools employ standardized interfaces, then tools have the potential to share context and data that can be correlated automatically before it is shown to analysts. Security standards also allow analysts to more easily develop custom tools that automate repetitive, manual tasks and allow analysts to spend more of their time and attention on challenges that require expert knowledge and human insight. Together, protocol, data, interface, security, and other types of standards limit attackers' options, make attacks easier to detect, and can improve and accelerate the enterprise response to attacks.

In theory, creating uniformity of practice also means that, if an attacker finds vulnerabilities in that practice, as codified by a standard, they can use that vulnerability against a diverse set of targets. In practice, unification of practice has been found to be far more beneficial to the defender for two reasons. First, it is relatively uncommon for vulnerabilities to exist in a technical standard itself and far more common for them to manifest in specific implementations of that standard. Because of this, attackers' efforts usually need to be directed at individual applications rather than all implementations of a standard. Secondly, unlike physical terrain where the concentration of defenses can sometimes be nullified by having an attacker circumnavigate them, in cyberspace the only vectors available to an attacker are those that are created by the defender's infrastructure. If a given protocol or data format is not used by any application or device in a defender's enterprise, there is no opportunity for an attacker to use that protocol or format to compromise the enterprise. In effect, the attacker is dependent on the defender to provide the attack paths, such as they are, into the defender's enterprise. Given this premise, it is clearly to the defender's advantage to limit the number of attack paths and to concentrate defenses along those attack paths that do exist. Technical standards facilitate this by allowing multiple devices to use a common protocol or data format, effectively reducing the number of attack paths into the enterprise and allowing a common point of defense. For these reasons, the uniformity of practice facilitated by technical standards is deemed strongly beneficial for defenders.

Of course, technical standards alone are not enough to defend connected devices and infrastructure. Limiting an attacker to a small number of attack vectors does little good if those vectors are undefended and the targets are open to attack. For this reason, standardization must be accompanied by other protections at the network, data, and application layers of the enterprise. However, such protections are easier to design, as well as more commercially viable and thus of greater interest to vendors, if these protections can focus on a small number of attack vectors with broad applicability. Since standardization of practice will reduce attack vectors and broaden adoption of a single practice, it is beneficial for such standards to be created sooner rather than later to give focus to the development of appropriate defenses. The bottom line is that standardization of practice will often need to lead development of defensive technologies.

3 US Government Role

Having noted the important role technical standards play in the securing of IT devices and the networks to which they connect, what should the US government's role be in the development of such standards? In the modern US, standards development efforts tend to be dominated by industry and most are created "with little or no government involvement". (Chopra, Sapiro and Sunstein 2012) Given that observation, it is reasonable to question whether it makes sense to deviate from this general practice and have the US government be more active in the generation of IT standards.

In fact, there are multiple reasons for the US government to play an active role in the development of IT standards. The following sections look at some of the most important reasons for government involvement in the creation of IT standards where those standards play a role in cybersecurity capabilities.

3.1 Expertise

The US government possesses significant IT expertise. Multiple groups within the US government, including but not limited to the DHS Science and Technology Directorate (US Department of Homeland Security 2015), the NSA/CSS Research Directorate (US National Security Agency/Central Security Service 2015), and the Computer Security Division of NIST's Information Technology Laboratory (US National Institute of Standards and Technology 2015), all perform cutting-edge IT research that could help inform standardization efforts. This general expertise and the forward outlook of these researchers can help in the creation of standards that help advance IT and cybersecurity capabilities.

Beyond this technical expertise, the US government, has significant prior experience with the creation of IT and cybersecurity standards. The list of IT standards in which the US government has played an important role is very large, starting with the Defense Advanced Research Projects Agency's (DARPA) work on the original ARPANET protocols. (Defense Advanced Research Projects Agency n.d.) Some examples of recent standards in which the US government has had an important role include, but are not limited to, CVE[1], TPM[2], OVAL[3], CWE[4], XCCDF[5], STIX[6], CYBEX[7], TAXII[8], and SWID tags[9]. Many of these efforts have been adopted by formal standards bodies, although some represent “de facto” standards – broadly used but lacking formal adoption by a recognized standards body. For reasons elucidated in section 5.1, virtually no standard is published without its detractors and this paper is not asserting that any of these efforts were perfect. However, the US government’s familiarity with the process of technical standard creation make it one of the most experienced organizations in IT.

3.2 An Important Stakeholder

The US government has a direct interest in the overall security of US cyber infrastructure, both in private industry and in the government itself. Cyber attacks, especially those directed against major economic activities or critical infrastructure, could be severely disruptive to the nation. The Office of the President of the United States notes that, “in limited policy areas where a national priority has been identified ...a convening role by the Federal Government may be needed to accelerate standards development...” (Chopra, Sapiro and Sunstein 2012) Inadequately secured IT devices can be both targets of attacks and vectors through which other connected systems are attacked. Consequently, the US government has a strong interest in seeing that security solutions for IT, including the development of new IT and cybersecurity standards, are initiated.

The US government has the responsibility to protect US citizens and interest from threats both domestic and foreign. This includes threats that operate in cyber space. To address these challenges, the US government undertakes efforts that include, but are not limited to, improving cyber defenses, coordinating responses to attacks, and identifying new cyber threats to US interests. In the course of performing these duties, the US government engages with an array of organizations in order to better understand their security needs. Within the US, this broad, cross sector perspective is arguably unique to the US government. As such, this perspective constitutes an invaluable asset to standards development efforts that seek to create a broadly usable standard. This means that the US government not only has strong interests in the cybersecurity of the country in its own right, but insight into the interests of others as well.

[1] CVE = Common Vulnerabilities and Exposures (https://cve.mitre.org/index.html)

[2] TPM = Trusted Platform Module (http://www.trustedcomputinggroup.org/work-groups/trusted-platform-module/)

[3] OVAL = Open Vulnerability and Assessment Language (https://oval.cisecurity.org/)

[4] CWE = Common Weakness Enumeration (https://cwe.mitre.org/)

[5] XCCDF = eXtensible Configuration Checklist Description Format (http://www.iso.org/iso/iso_catalogue/catalogue_tc/catalogue_detail.htm?csnumber=61713)

[6] STIX = Structured Threat Information Expression (https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=cti-stix)

[7] CYBEX = Cybersecurity Information Exchange Techniques (http://www.itu.int/en/ITU-T/studygroups/com17/Pages/cybex.aspx)

[8] TAXII = Trusted Automated eXchange of Indicator Information (https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=cti-taxii)

[9] SWID tags = Software Identification tags (http://www.iso.org/iso/catalogue_detail?csnumber=65666)

3.3 An Important User of IT

In addition to its interest in the overall cybersecurity of the nation, the US government is a direct user of IT products to manage and defend its own networks. As a user of IT products, the US government is invested not just in the ends of achieving improved IT security (as discussed in its role as a stakeholder), but in operational aspects of the solutions available to and their compatibility with US government internal requirements. With regard to these requirements, government interest in certain standards serves as a signal to vendors who are deciding whether support of a given standard in their products is a worthwhile financial investment. For these reasons, the US government not only brings a valuable perspective to the design of a standard (reflecting their operational requirements), but can increase vendor interest in the standard simply by being seen as a likely customer for tools that adopt that standard. This latter role should not be overstated – for major vendors the US government constitutes a small percentage of their customer base, and the government market is just a tiny fraction of the overall IT sector. Nonetheless, customer interest is an important consideration for vendors in standard adoption, and the US government is a sizable customer.

Beyond protection of its own systems, the US government is actively involved in the protection of important non-government systems, such as its efforts to protect critical infrastructure. These too can influence IT usage in these industries. For example, the “Risk Based Performance Standards Guide” (US Department of Homeland Security 2009) includes cybersecurity requirements that apply to all “high-risk chemical facilities”. While this guide does not name individual technical standards, US government involvement in a standardization effort related to one or more of these requirements would be of interest to chemical facilities seeking to ensure compliance with US government guidelines. As a result, the government’s role as the source of regulations and requirements that apply to certain markets beyond the government means that government interest in specific IT standards can be seen as carrying extra significance.

4 Technical Standards - An Overview

A technical standard attempts to create uniformity in a specific activity employed by a particular community of practice. Some standards are “de facto”, meaning that they are broadly adopted in products without being officially recognized by a standards body. These stand in contrast to “formal” standards which do have the endorsement of a recognized standards body (although the definition of “recognized standards body” is not universally agreed upon). While these are both forms of technical standardization that implementers can use in designing their products, for many implementers and users the distinction is important and almost always comes in the form of a preference for formal standardization. It is important to note that de facto standards do have value to users and implementers, can see wide adoption, and (as discussed in section 5.2) can be a more effective route to the eventual creation of a formal standard than starting that effort within a standards body.

A technical standard itself is a specialized form of document that has its own conventions and rules. Its goal is to describe, usually in prose (rather than, for example, through computer code), a set of technical behaviors in such a way that all readers come to the same understanding of those behaviors using only the standard itself as a guide. This is a daunting challenge. As the Scottish philosopher Thomas Reid observed, “There is no greater impediment to the advancement of knowledge than the ambiguity of words." To compensate, standards developers need to be both precise and thorough in their descriptions. Specialized expertise by experienced standards authors, as well as the rigor and review imposed by the processes of most standards bodies, are necessary (and sometimes, even then, not completely sufficient as discussed in section 5.5) to produce a document that can create consistent, uniform practice among all its readers.

4.1 Types of Technical Standards

Technical standards support a range of activities and roles, and different types of standards attempt to support different types of outcomes within the communities they support. Examples of standards types include:

Interface standards - These define how two components are expected to interact with each other. The goal of such standards is to allow these components to interoperate in a specific activity using only the standard as a guide. Examples of this type of standard include Internet Protocol v4 (IPv4, RFC 791), Transmission Control Protocol (TCP, RFC 793), and as well as many Application Programming Interfaces (APIs).

Data format standards - These define how particular types of information are represented. The goal of such standards is to allow common understanding of the structure and meaning of this data. Examples of this type of standard include the number encoding portion of the IEEE Standard for Floating-Point Arithmetic (IEEE 754-208), ITU's X.509 certificate, and the Comma Separated Value format (CSV, RFC 4180).

Vocabulary standards – These provide standardized names or "controlled vocabularies" within a particular topic. The goal of these standards is to provide a set of terms with agreed-upon meanings. Examples include country codes (ISO 3166-1); language codes (ISO 639); the ASCII (ANSI X3.4-1986), Unicode (ISO/IEC 10646:2014), and other character sets; and Common Vulnerabilities and Exposures (CVE) names.

Process standards - These define how to perform a specific activity. The goal of such standards is to ensure that different parties reach the same results when given the same inputs. Examples of this type of standard include the operation handling procedures of the IEEE Standard for Floating-Point Arithmetic (IEEE 754-208), hash algorithms such as SHA256, and encryption algorithms such as 3DES.

Framework standards - These define roles and relationships within multi-component architectures. The goal of these standards is to create a common understanding of roles within a larger framework in order to help vendors develop products that align with those roles, and thus the framework itself. Examples of this type of standard include the Open Group's "Dependability through Assuredness Framework" (O-DA), the Trusted Computing Group's "Trusted Network Connect Architecture for Interoperability", and NIST's "Framework for Improving Critical Infrastructure Cybersecurity".

Note that these are general categories and an individual standard could encompass aspects of multiple types of standards. For example, IEEE's Standard for Floating-Point Arithmetic defines both a data format and a process. There are other ways one might categorize standard types, but the important point with regards to this paper is not the categories themselves, but the observation that different standards can play different roles in unifying practices. These different roles mean that the different types of standards require different types of expertise in their creation, and different criteria need to be used to judge their effectiveness. For example, a standard's ability to facilitate immediate interoperability among products is a critical criterion for success for interface and data format standards, but is less applicable to more high-level framework standards.

In summary, "technical standards" are a relatively broad genre of technical literature and, as with all literary genre, there can be disagreement on the criteria for inclusion. For example, some argue that only recognized standards bodies create technical standards, thus excluding de facto standards and publications by parties other than those bodies. This paper is primarily concerned with standards that are created within formal, public organizations, but does not limit itself to organizations that are universally recognized as "standards bodies". Similarly, different types of technical standards have different objectives. However, all technical standards share a need for rigor and precision in their contents if they are to have the desired impact of unifying practice among disparate parties.

4.2 Standards and New Technologies

There are many existing standards developed for traditional endpoints such as workstations and servers. Many of these standards can be, and are, employed by new technologies. Standards like XML, HTTP, and TCP/IP networking are employed in numerous new technologies such as medical devices, IoT, and similar. However, new technologies can introduce new domains of operation that have not previously undergone standardization, such as new types of information structures. Even when there are standards that have similar roles, often new technologies have different requirements and priorities than traditional endpoints, and these differences can make the use of existing standards impractical. For example, IoT devices might have limited processing capacity or limited memory (at least, when compared with traditional endpoints). They might have limited bandwidth or lack persistent network connectivity, and might need to be frugal with power consumption. In addition, some new technologies bring new operational constraints. For example, by their nature, most medical devices have a significant role in human safety. This impacts the level of failure tolerance that these devices can support. These factors or constraints were often not considered, or only partially considered, when standards were developed for traditional endpoints. As a result, while some existing standards can (and generally should) be used within new technologies, the specialized constraints and circumstances of these technologies can sometimes make those standards unusable without modification. Attempting to shoehorn an existing standard into a new situation it was not designed to handle can lead to poor performance or require so many proprietary extensions to the standard that the benefits of standardization are lost. As a result, existing standards might need to be modified or new standards created to deal with the different needs associate with new technologies.

5 Creating a Technical Standard

This section considers the process by which technical standards are created. It includes observations collected from individuals with extensive experience in the development of technical standards. Their observations include key factors that can impact the odds that a specific standardization effort will reach fruition, resulting in a useful and impactful technical standard. It should be no surprise that many factors influence the success of a standardization effort, and not all these factors are under the control of any one individual or group. Given this observation, it should be recognized that no prescriptive set of actions or behaviors necessarily guarantees the success or failure of a given standardization effort. That said, certain factors and decisions can increase the risk of failing. This section identifies several of these pitfalls.

5.1 A Successful Technical Standard

Most participants in a standards engineering effort want the effort and its resulting standard to succeed.[10] However, as observed by Carl Cargill, determining what constitutes success for a standard can be difficult. (C. F. Cargill 2011) Certainly adoption by vendors and impact on practice could be viewed as success criteria, but either can happen without the other.

Adoption without impact occurs when vendors adopt the standard in their products, but it ends up seeing little or no operational use. Impact without adoption is rarer, but can also occur. One example was the effort to create the Open Document Format (ODF). The effort was initiated by Sun Microsystems in an attempt to compete with Microsoft's proprietary file formats. While the ODF standard was ultimately endorsed as a formal standard by both the OASIS and ISO/IEC standards bodies, it saw little adoption. However, ODF was arguably one factor in Microsoft's move to create an open standard for its own file format, Open Office XML in ISO/IEC, which had a significant impact on available capabilities for implementers. (C. F. Cargill 2011)

It is also the case that two groups can look at the same standardization effort's outcome with a completely divergent assessment of that effort's success. A particular vendor might feel a standard fails to fit with their product models, while another vendor feels it fits well with theirs. End users can also reach different conclusions if a standard meets the operational needs of one party, but fails to meet the needs of another. A corollary of this observation is that two participants in a standardization effort, both firmly invested in the effort's success, can have diverging opinions of what success looks like.

The important lesson to take from this is that, when initiating participation in a technical standards creation effort, an organization should give thought to what its own understanding of a successful outcome looks like. This outcome needs to go beyond the mere publication of a standard and should include who they desire to be impacted by this standard and how that impact should manifest. This understanding of success might differ from other participants in the same effort, but it is important to note that differences do not necessarily mean these multiple understandings are in opposition. Recognizing that others might have different goals and priorities in the effort is useful to keep in mind when negotiating consensus on specific features and capabilities.

[10] On occasion, standardization efforts include members whose goal is to ensure that no standard gets created or that any created standard is ignored. This might occur in response to perceived competition between the standards effort and existing products or standards the member favors. Such opposition can be subtle and might take the form of reasonable suggestions that ultimately stall forward progress, create scope creep, misalign the work with its target market, or otherwise increase the effort's risk of failure. (C. F. Cargill 2011) The best response to such attempts is to have a leader who is able to maintain the effort's scope, timetable, and focus, and keep the work aligned with the interests of a majority of participants.

Recommendation: A participant in a standardization effort should have a clear understanding of what constitutes a successful outcome of their engagement in that effort, while recognizing that other participants might have different, but not necessarily incompatible, definitions of success.

5.2 The Sources of Standards

New standards can arise in a number of ways. This section looks at three of the most common sources for new technical standards.

Revision of Prior Standards

Some efforts seek to amend or expand an existing formal standard. This might be undertaken to update a standard to support new conditions that arose after its initial publication or it can help manage post-publication feature requests and error fixes in a way that does not split the standard along separate paths. Numerous standards have undergone revision during the course of their lifetimes including the Internet Protocol (IP), Network Time Protocol (NTP), Post Office Protocol (POP) and many others.

The advantage to this approach is that, because the process starts from an existing standard, there is already a written document to work from. There is also usually an existing community of interest to guide development, including both vendors and consumers. The disadvantage of starting from an existing standard is that often there is a strong desire for backwards compatibility, which can limit the amount of change possible and might force the revised standard to retain unwanted aspects. A “major revision” implicitly suggests removal of the need for backwards compatibility, but such drastic changes can be more disruptive to users of the existing standard, complicating adoption. For example, the transition from IPv4 to IPv6, which are not interoperable with each other, has been ongoing for over a decade, and will likely continue for an indeterminate time into the future. Until the transition is complete, vendors are sometimes forced to support both protocols to ensure network connectivity.

Codification of Existing Practice

In some cases, a standardization effort seeks to codify a set of practices that are already used in existing products but which have not undergone formal standardization. This can include turning a de facto standard into a formal standard, or could involve a vendor offering an influential, proprietary practice for standardization. This can help increase adoption and help to prevent multiple incompatible variations of the practice from arising due to a lack of a commonly recognized authority defining the standards of that practice. One example of codification of an existing practice was the standardization of the Ethernet protocol, which was commercially deployed years before it was standardized under IEEE 802.3.

The advantage to this origin for a standard is that, since the behaviors being standardized have already been deployed in products, the standard is often well tested and demonstrably practical. The disadvantage to this origin is that it often relies on a vendor’s willingness to share their proprietary work with a standards body. If the specification is the product of a single vendor, they may be reluctant to give up control and may oppose any deviations from their own practices. On the other hand, if there are already multiple vendors implementing slightly different versions of the de facto standard, unifying around a common model can be extremely difficult as each vendor has a strong financial interest in standardizing around their own variant.

Something like the latter situation happened when IEEE attempted to develop a single standard for Local Area Networks (LANs), but was faced with three different technologies each backed by a different, significant vendor. Ultimately, reconciliation proved impossible and IEEE was forced to create multiple LAN standards, each with their own niche use. (C. F. Cargill 2011)

Addressing Emerging Needs

Other times standards might arise because of an unmet need that can be filled through standardization of practice. The need could arise from a range of factors such as customers seeking greater transparency or interoperability in the tools they purchase or vendors seeking to create or expand markets. The advantage of this origin is that it does not require the presence of pre-existing work, nor are the authors beholden to align with prior work. In other words, the standards development team is able to work from a clean slate. However, this clean slate also leads to significant disadvantages in that it requires recruiting knowledgeable, invested parties to help design the standard. It requires these parties come to a consensus, not just around a solution that meets the noted need, but around what needs are to be met in the first place. It then requires market forces to encourage adoption of the resulting standard so that it has an impact. Development of a new standard from existing standards or de facto practices also requires these factors, but their challenge is usually just to coalesce existing groups, designs, and markets around the revised/codified standard. A standardization effort that starts tabula rasa must grow all of these interest groups from nothing. As a result, starting from only an observed need is arguably the riskiest beginning for a development effort.

It is for this reason that standards bodies are often not the best place to perform the initial design work of solution creation. While there are numerous examples of successful standards that did, in fact, begin within standards body working groups, it can often be more effective to bring a standards group an *initial solution* that has already undergone some degree of review rather than to convene that group around a *problem*. Starting a standardization effort with a proposed solution helps to concentrate the group, not only in development of a consensus solution, but in its understanding of the core problem to be solved and the goals of the solution. Many standards groups make modifications to a proposed solution. These include changes that attempt to modify goals or address different needs, but those changes have a better chance of being clearly articulated in the presence of an initial solution that acts as a baseline. Without such a baseline, groups can find it more difficult to come to a consensus regarding the underlying need or ultimate goal of the effort. Even worse, groups can operate for months but fail to recognize that they are operating without the benefit of such a consensus. For these reasons, even if the initial proposal does not meet the ultimate needs of the group to which it is submitted, it still helps accelerate consensus by providing a concrete subject for revision. Finally, it is worth noting that a common complaint against standards bodies is that “they take too long to create standards” and starting with a draft solution is an effective way to reduce the amount of time needed to produce a published standard. Ideally, development of an initial solution is done with a small group of knowledgeable, invested parties who can then support further refinement of this solution when it is offered to a standards body.

Recommendation: When possible, convene a workgroup within a standards body around a draft solution rather than around a problem.

Of course, when one already has an initial solution that meets most of one's own requirements, there can be a temptation to shop around for a standards body willing to "rubberstamp" this work without modification. Doing so is almost always counterproductive in the long run.

One important requirement for the generation of a successful, widely adopted standard is to have a diversity of perspectives represented in its design. Doing so helps to create a standard that meets the needs of a broad group of users, which incentivizes adoption by vendors. (Section 5.3 goes into more detail about the different perspectives that should be included when creating a technical standard.) In most instances, the creation of an initial solution, while a valuable first step in the standards creation process, does not include input from a sufficiently diverse set of stakeholders necessary to create a broadly useful standard. Attempting to convert an initial effort into a final standard without additional input risks creating a standard that few find useful and that fails to achieve broad adoption. It does make sense to "shop for" a standards group that is aligned with the overall vision and market interests under which the initial solution was developed (as is discussed in more detail in section 6.1). However, the objectives of engaging with a standards body should include using that body to increase the diversity of participation. These diverse viewpoints can then be employed to evolve the initial solution so that it can meet the needs of more than just a small subset of stakeholders.

Recommendation: Avoid having an initial solution "rubberstamped" by a standards body. Instead, use engagement with a standards body to evolve an initial solution so that it better meets the needs of a broad set of stakeholders.

It is important to note that, even when dealing with very new technologies, all three creation methods are possibilities when creating standards that support that technology. An existing standard that was designed for IT practices might undergo a revision that allows it to be more easily applied to the needs of new technologies. Ad-hoc and proprietary solutions employed by new products could be offered for standardization and evolved to meet the broader needs of a developing community of practice. And, of course, some new needs, especially those that are more unique to a new technology, might need to be addressed through the engineering of new solutions.

The important thing to remember is that just because the context of the standard is new does not mean that existing solutions, be they existing formal standards, de facto standards, or even proprietary, cannot serve as a basis for addressing this new context.

5.3 Important Roles in Standard Creation

When developing a technical standard, it is important to include a diversity of participants in the process. Participants need to represent a recognizable community of practice and thus have some uniformity in their perspectives on the standard's area of operation. However, having a broad group of active contributors within such a community helps to ensure that any products of the standardization effort reflect the needs and perspectives of, and thus are usable by, the community as a whole rather than just a subset thereof.

Any standards development effort requires knowledgeable, invested parties actively participating in the standard creation process. These participants must be knowledgeable in order to produce a standard that works and which reflects the realities of current needs and technologies. They must be invested in the outcome in order to ensure that the effort is focused on producing a result that they are willing to use rather than just indulging in an academic exercise. Finally, the effort requires active participants who will help develop drafts and provide specific, detailed feedback on those drafts.

While there are different ways to categorize participants in a standardization effort, this paper identifies four broad types and generalizes as to the role they play in designing a standard:

End Users - These are the groups and individuals that are trying to solve operational problems using the proposed standard. Most end users are primarily focused on the capabilities created in the standard and the degree to which those capabilities meet their operational needs. End users tend to get involved in standardization efforts because they have a need that cannot be adequately met with existing products and they believe standardization will address that need.

Implementers/Vendor - These are the groups that create the products that conform to the standard. Vendors tend to focus on practicality (sometimes voiced as "frugality") in the standard to ensure that it can be incorporated usefully into their products without making those products overly expensive. Vendors tend to get involved in standardization efforts because they feel that supporting the standard can create or enhance markets for them, or that the standard will reduce current production costs.

Technical Solution Experts - These are groups and individuals skilled in technical solution development. They might not necessarily be experts in the domain that the standard covers. This group helps develop technical solutions to reflect user desires and vendor concerns. In other words, they help synthesize the higher-level capabilities and constraints expressed by end users and vendors and translate those attributes into technical mechanisms and models.

Standards Experts - These are experts in the creation of standards documents in general, and in the procedures and methodologies of the standards body where the work is occurring in particular. Members of this group focus on the codification of agreed-upon technical solutions into a suitably expressive and precise document or documents. They also help the standards development team comply with any special requirements that the standards body itself might impose. Additionally, they help ensure that the result is acceptable to any oversight the standards body requires prior to formalizing a specification as a standard.

It is important to have a balance of these groups in any standards creation group in order to achieve optimal results. If vendors and implementers are excluded, the solution developed can easily become a capability wish list that is too impractical for adoption. A group of only technical solution experts can end up creating an elegant standard design that fails to meet real-world, operational needs. Both NISTIR 8074 (Hogan and Newton 2015) and OMB M-12-08 (Chopra, Sapiro and Sunstein 2012), while primarily concerned with the role of standardization to meet US government needs as standards users, recognize that these efforts cannot effectively occur without support from vendors as well. Moreover, because not all end users and vendors have the same requirements, and because no single technical solution expert or standards expert is capable of performing all corresponding work in a standards development effort, it is necessary for there to be multiple instances of each participant type. While different situations might call for different balances of each role, having multiple representatives from all roles increases the chance of creating a standard that meets real needs, is usable in real products, and which facilitates the level of common understanding that is ultimately the goal of any standardization effort.

Although it is uncommon for a single party to represent both an implementer/vendor and an end user of a single standard, a single party could embody virtually any other combination of participant types. Conversely, participants who represent only one of these types can still contribute valuable insight into a standards development process. In particular, some standards development groups manifest a subtle prejudice against participants who are not technical solution experts, reflecting a bias towards technical engineering skills and participants who can participate in solution development. However, collecting user and vendor requirements is critical in developing a standard that meets operational needs and can be practically employed. Inputs from multiple users and vendors, even those who are not highly technical, can help address these aspects of standards development.

The objectives of some participants can appear to be or actually be in conflict. For example, users might want a capability that vendors feel is not cost effective. Similarly, a design that is practical in one vendor's implementation might be deemed impractical in another vendor's implementation. While this can create conflicts, it is not necessarily a zero-sum situation. If a compromise can be reached that addresses the more critical of each parties' requirements, and partially benefits both participants, the result is often a more robust, resilient standard that meets the needs of a broader community of practice.

When developing standards in support of newer technologies, the relative immaturity of that technology makes the need for broad participation in developing standards especially important. Many aspects of traditional IT have developed norms of practice that have produced a level of uniformity. For newer technologies, similar norms of practice are less likely to have formed and those that have formed are likely to be less mature. This means that communities of practice formed around newer technologies often have less uniformity when compared to traditional IT communities. Capturing this greater diversity of current practice (and hopefully unifying it around a common set of activities) requires that this diversity be represented during the standards development process. This requires broader participation of end users and implementers than might be needed when developing traditional, better established IT standards. In addition, sometimes new technologies touch on sectors or practices with which traditional IT has had little contact. For example, some classes of IoT devices bridge between cyber and different classes of physical systems. As a result, when creating standards supporting such IoT devices, communities of users and implementers whose interests are primarily in the physical side of device operation might need to participate in the development of certain types of standards. For example, if one were creating a standard for communicating with network-connected door locks, it would be important to have someone familiar with building fire codes, physical vulnerabilities of the locking mechanisms, handicap access, and other physical considerations that might have an impact on the developing technical standard. This is a perspective that most existing computer standards development organizations likely do not normally include. Accordingly, it may require special efforts to bring such participants into the standardization process. In all cases, when standardizing around new technology, extra effort should be taken to ensure participation by representatives of the technology's whole community, rather than just the portion that is focused on the IT aspects of the technology.

It is worth emphasizing that the US government has the ability to represent the perspectives of all categories of participant. As noted earlier, the US government is a major customer of cyber and cybersecurity products and thus represents an influential user perspective. Moreover, because government agencies have channels of communications to certain industry sectors, especially within critical infrastructure, it is able to experience a broader perspective of user needs than virtually any other organization. The US government sometimes directly provides services to the private sector and different parts of the government, thus filling the role of a vendor within certain contexts. The US government also includes engineers who have extensive experience developing and evaluating technical solutions to some of today's most challenging cybersecurity problems. These technical solution experts would constitute a boon to any standards development effort. Finally, US government representatives with experience in standards development can serve as standards experts in support of a development effort. Moreover, US government representatives can fulfill these roles in a way that is vendor agnostic. While vendors can also provide technical solution and standards experts, they are generally perceived (often accurately) as having a bias towards their own company's solutions. For all these reasons, the US government and its representatives can play a unique and highly valuable role in the development of standards. This does not obviate the need for other participants representing end users, lest the result be a standard that meets only government needs and thus is unlikely to see adoption except in expensive, custom implementations for government buyers.

[11] Likewise, other technical solution and standards experts should be involved, both to avoid having the government end up paying for all the work of the standards group, and also to better capture other technical perspectives. This said, the US government remains a valuable contributor, and its significant investment and focus in the protection of the national cyber ecosystem, argues strongly for active government representation in standards creation.

Recommendation: In any standard development effort, attempt to recruit engaged participants representing a diverse set of end users, vendors/implementers, technical solution experts, and standards experts.

5.4 Technical Standard Design Considerations

Once a body of participants is convened, either as part of a standards body or as a precursor to the involvement of such a body, the technical work begins. This includes: collecting requirements from participants, identifying a mutually acceptable set of practices around which to unify, and engineering a technical solution that will become the basis of a technical standard. This section looks at this process and provides recommendations to increase the chances of successful outcomes.

[11] Some in the government have suggested the creation of standards built around government requirements as a way to avoid the need to pay for custom software, but this benefit only occurs if vendors adopt those standards in their regular commercial products. This only happens if the vendors see a financial reason for this adoption, which almost always is based on perceived customer demand. If a standard fails to address the needs of all but government users, regular users will not demand it and vendors will not adopt it. While it might be easier to write purchase requirements around a formal standard, if that standard has seen no adoption, this is just another way of requesting customized functionality, along with its attendant costs.

5.4.1 Sound Engineering Practices

Engineering a technical standard shares similarities with other engineering efforts, including sharing the benefits achieved by following sound engineering practices. This may seem obvious, but standards workgroups, even within respected standards development organizations, can see uneven adherence to such practices.

There are multiple factors that can contribute to an incomplete or flawed application of sound engineering practices in standards development efforts. One factor is that not everyone involved is necessarily an engineer who has been trained in such practices. As noted in section 5.3, end user and vendor representatives, who are not necessarily experienced solutions engineers, can make valuable contributions to an effort in terms of identifying user needs and market priorities. However, these participants might not necessarily be familiar with sound engineering practices or their importance to the standardization effort. Moreover, even the dedicated engineers within a standardization effort will ideally represent a diversity of perspectives, which often means they come from a variety of different cultures. These cultures can be corporate, organizational, or even national. These cultural differences are reflected in the different practices each participant understands as important to an engineering process. Some participating engineers might engage in certain practices only because they are required by their employer, and ignore them when left to their own devices. This means that while the benefits of sound engineering practices might be “apparent” to all, their use is not necessarily automatic or even considered obligatory.

Formal standards bodies usually include procedures to support sound engineering practices, but informal bodies (such as might be used to create initial solutions prior to engaging a standards body) might need to be more deliberate in identifying and supporting these practices. Even within established standards bodies, the degree to which sound engineering practices are followed often depends on the preferences of the workgroup’s leadership. Moreover, it is important that the group explicitly agree upon the required practices in order to avoid situations where different groups are following different practices to different degrees. If workgroup leadership does not initially suggest such practices, participants may need to propose such practices to the group at large.

Some engineering practices that of special interest to standardization efforts include, but are not limited to:

- Careful scoping of the problem and solution spaces to keep the work manageable and constrain scope creep. (See section 5.4.4 for more on the risks of over-featuring.)
- Setting achievable and measurable milestones for progress with defined time limits to help drive and maintain forward progress.
- Creating document trails of decisions to reduce revisiting of resolved issues and to facilitate reconsideration of a decision should new information arise.
- Recognizing the need to balance capability with practicality, and a willingness to compromise in order to achieve that balance.

Recommendation: Agree upon and ensure that sound engineering practices are observed in the development of technical standards.

5.4.2 Market Awareness

Another consideration that technical standards development shares with many other engineering efforts is a need to understand the market that the result (be that a new product or a new technical standard) will serve. In the same way that a great product can fail because it does not fit with market needs and expectations, a new technical standard can fail if it is too far out of alignment with market expectations. Carl Cargill, citing a 1990 study by Martin Weiss and Marvin Sirbu (Weiss and Sirbu 1990), observed that an understanding of a standard's market was one of two factors most associated with standard success (the other factor being clear willingness by participants to provide written contributions to the workgroup, as described in more detail in section 6.4). (C. F. Cargill 2011) For a (somewhat clichéd) example of this, consider 8 tracks and cassettes. In theory, 8 tracks had the ability to record sound with greater fidelity than cassette tapes due to their faster play speed. However, 8 tracks suffered from a host of other disadvantages that impacted the quality of their playback and made them less convenient for users. The market valued convenience and playback quality over recording fidelity and (for this and other reasons) cassettes displaced the 8 track.

New technologies represent a new "market" for standards and the special requirements of their markets need to be recognized, especially as they can differ from the needs of traditional IT systems. For example, while traditional protocols might favor fidelity and availability of a service, IoT device manufacturers might be far more interested in minimizing bandwidth and power consumption in a protocol's use. While fidelity, availability, low bandwidth, and low power consumption are all desirable capabilities, what constitutes an acceptable trade-off between these capabilities differs across markets. In fact, different segments of the broad category of IoT devices have different priorities - the capability priorities in a home refrigerator or dishwasher differ from those of industrial control system devices. As a result, it is important to have an understanding of a proposed standard's target market and where the market's values lie. Part of the goal of a standardization effort can include evolving this market practice, but design decisions must still align with the market's values. This understanding should be explicitly captured by the working group since it is directly related to questions of scope and feature requirements for the developing standard. Failure to do so can lead to unproductive arguments about features and requirements caused by participants' differing understandings of the target market. Unfortunately, such disagreements, focused as they are on technical details, are unlikely to resolve the underlying issue of identifying core market values, and thus are likely to be repeated. Understanding of the market can be difficult to gain, but that understanding can be assisted by ensuring that end users and vendors who are embedded in that market provide input into the design process.

Market awareness is typically an area where US government participants have a great deal to add, but where it also has great need of input from others. The US government itself represents a significant market in its own right, and its requirements carry significant weight in the development of a standard. Organizations which have regular engagement with important commercial enterprises, such as government agency engagement with critical infrastructure sectors, can also help to devise requirements that reflect the needs of those enterprises as well. However, US government requirements often differ in important ways from commercial industry requirements. As big as it sometimes seems, the US government generally represents merely a niche of the total market many vendors support. To have a standard that meets the need of the broad market, and thus is more likely to see adoption, it is critical to include input from non-government entities. This increases the chances that any standard produced meets the needs and values of both government and the private sector.

Recommendation: Understand the values of the standard’s target market and align design decisions and tradeoffs with those values.

5.4.3 Detail and Diversity

There is often a desire to “create the one standard that meets all needs”. However, attempting to create a standard that meets the needs of many different practices, each reflecting a different set of interests, priorities, and market values, can create a standard that all communities see as overloaded and poorly suited to their specific needs, assuming it is even possible to reach consensus around a standard in the first place.

In general, experience shows that that there is an inverse relationship between the diversity of a standard’s target communities of practice and the degree of detail that the standard can prescribe in data and procedure. (Mann, Shapiro and Bodea 2014) By “diversity of communities of practice” consider the following continuum: hand surgeons, all surgeons, all medical professionals, all parties interested in patient health (i.e., medical professionals, patients, insurance adjustors, etc.). All members in this continuum have an interest in the health information of a patient, but the nature of that interest is more uniform among hand surgeons than it is among all parties with an interest in patient health. When considering “degree of detail”, one can think of another continuum: database instances, schemas, report forms, general classifications, and nominal IDs. Database instances require a great deal of detail in field names, data types, lookup keys and pivots, and precise understandings of most field values in order to support their typical operations. At the other end of the spectrum are nominal IDs, which are assigned specific meaning but themselves encode little or no detail. Examples of nominal IDs include IP addresses, license plates, and product serial numbers.

Experience suggests that attempting to support a diverse set of practices while simultaneously providing highly detailed codification of process and data is intractable. (Mann, Shapiro and Bodea 2014) In the previous example, the types of information a hand surgeon needs to record differ from the information needed by an oral surgeon, to say nothing of the difference between those and the details of interest to an insurance adjuster. A data structure that simultaneously standardized all the information needed by hand surgeons, insurance adjusters, and patients for their individual activities would be bloated and unwieldy at best. Some types of information, such as a patient identifier or an agreed-upon name for a particular medical procedure, can be usable across the entire spectrum of those interested in patient health, but these are examples of nominal IDs and contain little or no encoded detail. (They might serve as a look-up key into more detailed information records, but those records are not inherent in the IDs themselves and the records could change without any alteration of the ID.)

The implication to be drawn is that when developing a standard to support a community that covers a diverse set of practices, care should be taken to focus development on activities that are shared across the community. While creating uniformity of practice where none existed before is a common goal of standardization, that uniformity needs to be created around activities in which all members of the served community engage, albeit potentially in different ways. If a standardization effort starts focusing on activities that are not of interest to a significant body of the community, this is a strong indication that the effort is attempting to include more detail in the standard than community diversity supports.

Of course, just because one has created a standard codifying a level of detail appropriate to the diversity of a community, doesn't mean that this standard is sufficient for all sub-communities therein. Some sub-communities may need to standardize their own practices to a greater degree of detail to meet their own specific needs. However, codifying around that particular set of practice should be an effort undertaken by that sub-community rather than within the larger community of practice. The sub-community might create a separate standard to meet their needs or codify specific use of extension points in the broader community standard to integrate their specific needs. (See section 5.4.5 for more on the important role of extension points.) In either case, it should be the sub-community that codifies a mechanism to meet their specific needs, rather than imposing those needs on a broader community that does not share them.

Recommendation: Avoid attempting to codify a high degree of detail in support of a highly diverse community. Detail in a standard should be balanced against the diversity in practice of the communities it serves.

5.4.4 Over-Featuring

Related to the previous point on detail vs. diversity, there can be a desire by engineering teams to identify and try to support multiple use cases in the standards they produce. The general reasoning is that, because different use cases have overlapping needs, it makes sense to add the features needed to support additional use cases rather than create a separate standard that has overlapping aspects or capabilities. This is often a reasonable and beneficial approach and can help bolster adoption by making the standard more capable. However, it is possible to take this too far and create a standard that becomes overloaded and difficult to use.

The typical way in which multiple use cases are supported by a standard is to require support for a common set of core features and then make features associated with individual uses optional. The reasoning is that users and vendors can support the use cases appropriate to their needs and business model and ignore the others. This can be a reasonable approach, but it needs to be understood that optional features are not zero-cost. Optional features complicate interoperability by allowing multiple "conformant" tools that have different capabilities. This complicates the lives of end users by making it more difficult to determine whether the tools they purchase are interoperable with regard to the optional features they want. This, in turn, can lead to a perception that the standard does not actually create interoperable capabilities. To combat this, some vendors feel compelled to support all optional features of the standard, but this increases the vendor's implementation costs - cost that are difficult to justify if some of the optional features are not aligned with their product's core capabilities. As a result, while including a large number of optional features in a standard can appear to be an efficient and low-impact way to make a standard more useful, these optional features can ultimately work against successful adoption and deployment of the standard.

Of course, if one makes support for all use cases mandatory this addresses the challenge of interoperability, but it also all but guarantees that both users and vendors end up supporting aspects of the standard that they do not need or want. The result is a greater cost to implementers (which is then passed on to customers) as well as a greater management cost to users who need to ensure that their enterprise can support all mandatory aspects of the standard.

In addition to creating too many features by trying to make a standard overly broad, one can also over-feature a standard by making the requirements overly deep. This involves creating normative requirements that are too closely bound to one way of supporting a given use case. Doing this makes the standard more complicated (due to the increased number of requirements) and makes it less flexible as well. An example of this could be a situation where a group standardizes a data representation, but then goes further and creates requirements about how that representation gets used. While the additional requirements might make sense for the targeted use case, they would preclude conformant implementations from meeting the use case in different ways, and could prevent support for other, valid use cases of the standard. Of course, in some cases, it might be appropriate to bind representation and use, so situations need to be examined on a case-by-case basis. Standard development teams should carefully consider each normative requirement to verify that it serves the central goal of the standard rather than codifying one of potentially multiple valid uses or implementations of the standard.

Standards development efforts should be conservative in the number of use cases, supported feature variations, and other requirements that a standard supports. In some cases it is highly beneficial to support closely aligned use cases and avoid the need for redundant standards. Likewise, the ability to support some level of variability through optional features is not just beneficial but necessary in some cases (as is discussed in the following section). Similarly, it can be necessary to tie a standard to certain details of use cases to ensure interoperability. However, care must be taken with regard to the overall requirement footprint of the resulting standard. It might be necessary to prune some use cases or use case details which, while reasonable on an individual level, cannot justify the additional cost they invariably bring to those seeking to adopt and/or use the standard.

Recommendation: Be conservative in the number of use cases and features (optional or required) a standard supports.

5.4.5 Extensibility

As has been stated before, the primary goal of a technical standard is to unify a particular practice within a specific community. However, the practice in question does not take place in a vacuum. While there might be agreement about a standard’s practice, vendors might want, and their customers might demand, additional features and capabilities beyond what is included in the standard. These additional capabilities might represent specialized needs of certain sub-communities or vendor’s efforts to distinguish their product from their competitors. A standard that precludes such tailored use can be viewed as more of an obstacle than a benefit by both vendors and end users. For reasons described earlier, attempting to codify all possible customizations of a standard in the standard itself creates a standard that is over-featured and difficult to use and adopt, and often such an enumeration of customizations cannot be known in advance.

The solution to this challenge is built-in extensibility. This can be accomplished in two different ways. For less detail-based standards, such as nominal IDs, or for more narrowly scoped standards where the specific practice it codifies has little variation (such as timestamps), customization is usually supported by ensuring that the standard can be used within a variety of contexts. For example, a timestamp is frequently embedded in other data structures or message formats. This is done by keeping the standard compact and relatively flat (as opposed to hierarchical) in its representation.

The other way to address the needs of customization is to build extensibility into the standard itself. This is generally the approach for more detailed or complicated standards, such as data records or protocols that are intended to support a diversity of uses (such as HTTP). This can be accomplished by providing dedicated fields that different vendors and applications can use to add their own additional data and instructions. Commonly, these fields incorporate a customizable identifier, whose format is constrained by the standard so as to try to reduce the chance of two vendors claiming identical identifiers, and a field value. Tools use the identifiers to determine which of the extension fields they can understand and correctly process. The header fields in HTTP messages are examples of this approach. The use of a type field to support an extensible list of payloads or data formats in some segment of a record or message can be viewed as a variation of this strategy where the identifier is recorded in a dedicated location (i.e., the type field) and the payload is provided in another dedicated location. HTTP messages again provide an example of this in their Content-Type field, which identifies the data type of the message body.

When designing extensibility into a standard, it is important to ensure that this extensibility does not compromise interoperability. Typically, this is accomplished by including instructions as to how tools are to handle customizations that are not recognized, and including requirements that prevent dependency on such customizations for correct processing. In many cases, tools that receive a message or process a data record ignore customized fields they do not recognize. As part of this process, it is generally considered to be a good practice to explicitly constrain extension by limiting where customization can occur, and requiring that the presence of customization be clearly identifiable, even to those that do not understand or support a particular vendor or product’s customization. Incorrectly designed extensibility can lead to interoperability problems that compromise the utility of the standard, so it is vital that points of extension be carefully designed and their use tightly controlled.

One of the side benefits of incorporating extensibility is that it makes it easier to revise the standard in a backwards compatible way. If the standard needs to be revised, there will be a transition period, sometimes for many years, when both the old and the new version of the standard are in use in operational products. During this time, tools often need to support both versions of the standard in order to ensure interoperability. If revisions can make use of existing points of extensibility, effectively treating the revision as a codification of specific extensions within the initial version of the standard, then tools that only support the initial version of the standard will be able to operate without error when presented with a data structure or message from the revised standard, even if they are unable to support the new feature. This makes the revised standard much easier and less expensive for vendors to adopt, which encourages said adoption. By allowing extensibility in a way that makes backwards compatibility easier to maintain when the standard is revised, the standard’s capabilities can be expanded (and errors corrected) while naturally accommodating existing product deployments.

Recommendation: Understand how a standard supports customized use by vendors and users. If appropriate, support customization via well-defined extension points and procedures.

5.4.6 Quick Wins

Technical standards range in scope from small, narrowly scoped efforts, to architectures covering multiple roles and activities. There are successful examples of standardization efforts across the full spectrum of scope sizes. However, efforts that are both broad in scope and which attempt to codify a high-detail a high degree of detail across that scope are at a greater risk of failure than their simpler counterparts.

A distinction needs to be drawn between broadly scoped standards that are written at a high, abstract level, and broad standards that are written at a low, detailed level. For example, many framework and architecture standards focus on defining roles and high-level activities for multiple components in a system. Such standards can be largely abstract when it comes to the details of how certain interactions occur. However, a similarly broad effort that attempted to define the technical interfaces and data structures used in such an architecture requires far greater detail to realize. It is the latter that represents a more significant challenge.

This first reason for this is the inherent engineering challenge of developing a detailed solution for a large, complex system. It is almost invariably easier to engineer a solution to a small, narrowly scoped problem then to create a solution to a larger, more complex set of problems. This impacts both the chances of eventually creating a solution that is accepted, and also the chance that the solution has a flaw that is only discovered after the standard is published.

In addition, a large scope that touches on multiple actors usually means that it comes into contact with multiple communities of practice. While these communities might be aligned with each other in general, the communities can differ in their understanding of low level details, both in terms of what details are needed, and even in the meaning and role of “common” details. (The challenge of reconciling detailed information among differing communities of practice was discussed at greater length in section 5.4.3.) All of this can make it difficult to develop consensus on any solution, and can drive the group towards solutions that attempt to meet too many needs and thus become laden with numerous features that most sub-communities do not need.

Finally, complicated, multi-part standards face another challenge in getting adopted. Both vendors and users can become overwhelmed when faced with a large, complicated architecture or procedure. Customers might fixate on the degree to which their existing infrastructure needs to change in order to comply with the practices of the standards. Vendors whose products have significant overlap in scope with the standard might feel that the standard competes with their existing market and oppose the standard in the marketplace. On the other hand, a vendor whose products only overlap with a relatively narrow slice of the described architecture will only realize the benefits of adopting their portion of the standard if other vendors’ products also adopt corresponding parts of the standard - a factor over which they have little control. The result is that it can become harder to convince both vendors and their customers that the new standard is worth the effort to adopt and deploy.

Instead of attempting to standardize details for the totality of a large, multi-part procedure via a top-down effort, it is often more effective to take a bottom-up approach. In the latter approach, efforts focus on developing detailed standards for specific, narrowly scoped problems commonly seen within the larger overall vision. For the reasons outlined above, these more narrowly scoped standards are usually easier to design, easier to build consensus around, and easier to convince both users and vendors to adopt. Later, after some of these more narrowly scoped standards are published and see adoption, additional standards that define how to integrate the operations of those smaller, detailed standards can be built, thus enabling standardization of the larger, multi-part procedure. This "integrating standard" not only becomes easier to build since it is based on existing capabilities, but since both vendors and users have seen value in the component standards and might have already adopted some of them, the idea of the larger vision itself becomes a more attractive proposition. In addition, because the more narrowly scoped standards are created first and must be usable independently, one avoids creating unnecessary interdependence between components in the larger vision. All of this makes it easier not just to create useful component standards, but ultimately to realize the larger vision that drove them.

With regard to newer technologies, there is another reason to favor smaller standards. With newer technologies there can be a great deal of variation in how those technologies are used in an enterprise and marketed by vendors. Simpler standards often can support a wider diversity of deployment models and usage scenarios than ones that require the presence of multiple cooperating roles. For example, the timestamp format defined in RFC 3339 is used in a vast array of activities and operations. Had the authors created the standard with explicit linkages to some larger vision, such as adding features associated specifically with network traffic or with computer file systems, those linkages would have been a hindrance to this broad utilization of the standard. As the use of newer technologies matures and norms of use are developed, "common models" can develop and be standardized. In the near-term[12], for newer technologies it might make sense to focus on more narrowly scoped standards that that are better able to provide benefit across a diverse set of uses.

Of course, in order to identify narrowly scoped standards that can be readily integrated into a larger vision, it is useful to have a codification of the larger vision as a guide. Architecture standards written at a higher level of abstraction can be useful to identify potential component standards and can help identify points of connection those standards should support. However, it is recommended that those component standards be able to stand alone and provide utility without any dependency on the rest of the vision. Moreover, designers should be very wary of features that, while in line with the larger vision, provide little direct value to the specific capabilities the component standard is intended to support. Either of these situations are indicators of scope creep for the component standard that can complicate its development and damage its chances for adoption. Use the larger vision to identify needed parts and their roles, but be careful that development of those parts does not morph into an attempt to develop technical capabilities for the whole vision.

Recommendation: Smaller, more narrowly scoped standards have a greater chance of successful completion and adoption than large, multi-part standards. If possible, start with the former and, if desired, build the latter through the composition of those smaller component standards rather than attempting to engineer directly to the larger vision.

[12] "Near-term" in technical standards design is on the order of magnitude of about 5 years.

5.5 Follow-Up

One of the more common misperceptions about standardization is that, once the standard is published, the work is done. Many people incorrectly believe that once the standard is written and published, that adoption of and conformance with that standard will handle itself. Some go so far as to believe that the degree of adoption and conformance of a standard directly correlate to the quality of the standard and that attempts to encourage adoption and conformance interfere with "natural selection" that allows better standards to thrive while less effective standards are naturally displaced. This is not the case and actions taken after publication can mean the difference between a standard receiving widespread adoption or fading into insignificance.

5.5.1 Marketing

Standards require marketing and support in much the same way any other product does. Potential end users and product vendors need to be made aware of its existence and "sold" on its benefit to them. Customers need to see the capabilities and other benefits the standard creates for them. Vendors need to be convinced that use of the standard protects or increases their sales. The former often facilitates the latter as vendors frequently look at customer demand to determine whether support for a given standard is a good investment. Ultimately, vendors need to see some sort of commercial benefit for them to implement a standard, be that through the generation of new markets, the ability to market to a demonstrated community of customers who are demanding the feature, or simply to avoid losing customers to competitors who implement a standard customers see as beneficial.

Different messages are necessary for these different communities. End users are unlikely to be convinced by technical descriptions of the standard and are generally more receptive to demonstrations of how the standard solves their problems. Vendors are likely to respond to demonstrable customer demand as well as evidence that support for the standard does not require a major investment or production changes on their part. The bottom line is that the probability of any standard being adopted is increased by active outreach to build awareness of the standard and to market the benefits of the standard to its could-be users and adopters.

5.5.2 Interoperability

It is an unfortunate truth that, for all but the simplest of standards, it is relatively easy for two parties to comply with the standard without having completely unified understandings of the standard's operation. Cargill notes that this is a common problem in technical standardization efforts. (C. F. Cargill 2011) For interfaces and data formats, this can cause products to fail to interoperate under certain circumstances. For processes, this can create edge cases where the same initial conditions lead to different results.

To address this problem, it can be strongly beneficial to support efforts to encourage vendors not merely to be conformant with the standard, but to be interoperable with each other. (Hogan and Newton 2015) Such efforts are extremely important when a standard is first published. Thereafter, they become gradually less critical due to the existing body of interoperable products against which a new implementation can be tested. The Wi-Fi Alliance (Wi-Fi Alliance n.d.) was formed to address this particular need in wireless networking protocols. Early products supporting wireless networking suffered from frequent incompatibility problems. A group of wireless product vendors formed the Wireless Ethernet Compatibility Alliance (WECA), which later became the Wi-Fi Alliance, in 1999 to test and certify interoperability of their products. Today there is far greater interoperability of wireless networking products, even among those that do not have formal Wi-Fi Alliance certification. In most cases, it is not necessary to set up something with the level of organization created by the Wi-Fi Alliance. Hosting some "interop days" under non-disclosure agreements, where vendors can test interoperability between their products and address incompatibilities without fear of generating bad publicity, can be a simple, inexpensive way of facilitating interoperable standard adoption. Some standards development organizations support such events as part of their normal operations, but other organizations rely on outside groups to fill this need.

5.5.3 Reference Implementations and Documentation

Another way to facilitate interoperability and adoption in general is the production of reference implementations and similar tools that vendors can use to bootstrap and test development efforts. These tools do not need to be sophisticated, and bare-bones implementations might even be preferable so as to avoid the perception of competing with vendor products. These reference tools help clarify the expected behavior described by the standard. If these tools are released under appropriate licenses, the tools themselves or components thereof can be directly integrated into vendor products, reducing the investment cost, and therefore the risk, of adoption and accelerating adoption efforts. For standards that rely on multiple parties, such as communications protocols between two different roles, having a reference implementation for both communicants can be a big help to implementers by giving them a correspondent against which to test their own implementation. As such, simple implementation efforts can have a significant role in encouraging interoperability and adoption.

Additionally, following the release of a new standard with supporting documentation can help boost adoption. This documentation can help clarify questions that frequently arise when a new standard is first integrated into products and procedures. Documentation can also elaborate on expected practices and use cases that, while generally not codified by normative requirements, help to clarify the role the standard's authors expect the standard to play. This, in turn, helps implementers better plan their adoption strategies. Finally, publication of supporting material, especially in the first year or so after a standard is published, can help reassure adopters that they are not on their own and that there is an active community of support. Encouraging the perception that the standard has a living community behind it can make it far more attractive to would-be adopters.

It should be noted that there are standards that are widely adopted, implemented, and interoperable, yet never enjoyed any formal effort to market them or to facilitate interoperability among products. That said, standards that are supported by such efforts have a better chance of being adopted and usefully deployed in tools. A marketing effort cannot compensate for serious flaws in a standard. However, a flawless standard can still be ignored or fail to deliver interoperable capability if it remains obscure within its target community. For this reason, it is advisable that groups that invest in the time and effort to create a new technical standard protect this investment by allocating time and funding to support the new standard after its publication, at least until achieving a critical mass of interoperable adopters.

Recommendation: Be prepared to support efforts to raise awareness, advocate for adoption, facilitate interoperability, develop reference implementations and other resources, and otherwise provide support of a new standard for a period of time after the standard is published.

6 Working in Standards Bodies

A standard from a recognized standards body, especially one with an international scope, is usually more attractive to both potential implementers and to end users. Often end user organizations are more willing to set purchase requirements around formal standards than de facto standards. Similarly, a vendor with an international market is usually more interested in a technical standard from an international standards body since the standard is recognized in international markets and can lead to sales there. There is also often a perception (not always fair) that when compared to de facto standards, standards published by standards bodies have undergone a more stringent quality review, reflect the needs of a broader user community, and are more likely to be managed in the best interests of all members of the user community. For all of these reasons, endorsement by a recognized international standards body gives a technical standard a boost in its chance of broad adoption.

As noted in section 5.1, a standards body is not always the best place to begin designing a standard. That said, for the reasons noted above, those seeking to create broad adoption of a standardization effort might wish to eventually engage with such a body to create a formally recognized standard. Standards bodies represent their own challenges above and beyond the technical challenges of developing the standard itself. This section looks at some aspects of standards bodies and provides general guidance on how to work within them. The recommendations in this section can increase the likelihood that engagement with a given standards body produces a satisfactory outcome.

6.1 The Anatomy of a Standards Body

Some of the main international standards bodies have been in existence for many decades, but new industry consortia appear on a regular basis. (C. Cargill 2011) While all of these groups are, to some extent, open and include the development of standards among their activities, each of these groups is ultimately unique. Below are several ways in which standards bodies tend to differ from each other.

Who can be a member? - Typical options include individuals, companies or similar institutions (e.g., government agencies, academic institutions, etc.), nations, or some combination of these. In some cases, membership in the group is not open to those outside of a specific organization, but non-members can provide feedback and sometimes assist in the writing of standards.

How do you become a member? - Typical options include simply showing up (i.e., no membership fee), paying a membership fee, nomination by existing members, or appointment by some other organization.

How is consensus on a design decision measured? - Typical options include formal votes or determination by the working group lead.

Who can participate in formal votes (if any)? - Full voting rights might be given to all members or voting rights might be limited. In the latter case, it might be that voting rights are given only to members who have purchased a certain tier of membership, only to participants who have participated in some number of prior meetings, or some combination of these criteria. In some cases every eligible participant can case a vote, while in other groups each member institution only gets one vote regardless of its number of participants.

What is the process for adoption of a standard? - Differences include if oversight bodies must sign off on a standard, how many such sign-offs are needed, and how much time is allotted prior to each sign-off. One of the implications of this is the amount of time a body takes to go from a “finished” proposal to a published standard.

How are finished standards published? - Typical options include making the standard freely available or selling the text of the standard but providing a free license to all would-be implementers. Some standards bodies allow standards that include the use of proprietary technologies, meaning that adopters might need to pay license fees to some company in order to implement the standard.

What is the group’s volume of publication? - Some groups publish dozens of standards each year, while other groups might only publish one or two.

How are published standards received by users/vendors/other standards bodies? - The publications of some standards groups are universally recognized as authoritative. Others might only be recognized as authoritative within a nation, within a particular community of practice, or some other subset of the world. In some cases, one’s own organization might have a list of standards bodies that are viewed as preferable.

What is the typical member composition? - Some groups draw members exclusively from certain interest groups, while others draw from a broader set of communities. Some groups receive more participation from the open source community, from national governments, from specific global regions, from academia, or from other classes of participants.

How does the group “meet” to develop standards? – Groups can do most of their work over ongoing email discussions, periodic telephone calls, in-person meetings, or other methods. Even when groups support multiple methods, some will remain more important than others.

What are the typical types of standards created? - Groups might produce primarily standards for interfaces, data formats, processes, frameworks, or some combination thereof. Groups might also focus on certain types of technology, such as hardware, networking, applications, or other classes of technology.

What is the culture of the group? - Groups can be formal or informal, can be collegial or adversarial in their discussions, can be highly technical or highly political, or have other cultural characteristics.

These are some, but by no means all, of the ways in which standards development groups differ from each other. This large amount of variation between standards groups is an important factor for consideration by those seeking to create or formally adopt a proposed standard. Bringing a proposal to the wrong standards body can doom an effort to failure from the start. For this reason, it is important to evaluate a potential standards development group against both the needs of the standard development effort itself as well as the submitter’s own goals for the development effort.

Recommendation: Before proposing a new standardization effort in a standards body or consortia, evaluate the group to verify, among other things, that the group is a good fit for the proposed work, that the group can bring appropriate participants to the process, that the publication of the proposed standard by the group will impact the desired audiences, and that the procedures and culture of the group are acceptable.

6.2 Participant Reputation

The reputation of a given participant can play a significant role how their proposals are received, as well as their ability to recruit other participants to assist in writing, editing, and prototyping efforts. A participant who has a reputation for technical competence and/or an ability to find compromises and solutions that meet diverse needs usually finds more parties willing to give their ideas consideration and support. By contrast, those who show a lack of technical understanding, an unwillingness to listen to other ideas, or a constant bias towards a particular vendor and technology regardless of counterarguments, can find it more difficult to find support for their proposals and activities. One corollary of this is that long-term participants in a body generally enjoy a greater degree of support due to having an established reputation within the group. (Those who have bad reputations generally do not remain as long-term participants.) As NISTIR 8074 observes, "...long-term participation of the same USG representatives within an SDO[13] establishes trust and builds the credibility of those representatives. This is critical for effective communication and information-sharing and ultimately will assist in advancing the USG's strategic objectives in each SDO."

Similarly, there is often a quid-pro-quo at work in standardization efforts. For example, one participant might agree to review or edit a specification that is only of peripheral interest to them. Later, that specification's author can return the favor by assisting with work important to the first participant. This concept of social credit within standards bodies should not be underestimated, and investing time and effort to build such credit can pay off when one later needs others to support one's own high-priority efforts. Again, long term participation in a standards body, rather than showing up for specific efforts and then disappearing when they end, helps a participant build social credit that can help when the participant has a particular effort that is important to them. The combination of positive reputation and social credit can improve the chance of a participant receiving the support necessary for their priority efforts to be completed in a timely manner.

Reputation and social credit are more important in some groups than others. In general, it tends to be a larger factor when participants are seen as having some degree of independence in their activities. This is due both to the greater association of the individual with their contributions, and because participants have greater discretion to pay back any social credit earned. By contrast, in organizations where participants are explicitly just representatives of larger organizations, social credit tends to play a smaller role. For example, in organizations such as ISO and the ITU, where participants are representatives of nations and generally just represent decisions already made within national bodies, the reputation of individual participants tends to play a smaller role. Would-be participants in a body should look at the practices of that body to

[13] Standards Development Organization

evaluate whether investment in building social capital in the group is justified. For groups where participants are seen as having individual discretion in their activities (which includes many organizations where members participate on the behalf of companies or other organizations), the relatively small investment needed to support regular meeting participation and occasional editing of documents can have a payoff when the participant (or their organization) has a strong interest in an effort or topic.

It is important to emphasize that the reputation of a participant is not the same as the reputation of the organization that they represent. In most situations, individual reputation plays a larger role. Even if a particular organization is held in high regard, any benefit the participant receives from this association is likely to be probationary and can easily be lost. By contrast, a participant from an organization with a negative reputation (such as a company with a history of hijacking standards with proprietary, incompatible extensions) can have an uphill battle to gain goodwill in a group, but this can also be done. The US government always brings a reputation with it, either positive or negative depending on the predilections of the group they join. In either case, participants representing US government interests need to actively cultivate a positive reputation in that group as any positive associations can be fleeting, and any negative associations need to be overcome.

Recommendation: For standards bodies in which participant reputation and social credit play a role, actively seek to build this credit and consider supporting long-term participation in that body as a way to improve the participant's effectiveness in supporting efforts of importance.

6.3 Diplomacy

NIST's guidance on engaging in standards bodies (NISTIR 8074) observes: "Effective negotiation in standards development requires not just technical expertise by Federal agency participants, but a thorough knowledge of an SDO's standards development process and policies, as well as soft skills in negotiating with stakeholders with a range of often diverse and conflicting positions." (Hogan and Newton 2015) As noted before, participants can have differing, and sometimes even opposing, goals with regard to the objectives of a standardization effort. When one proposes a particular feature, solution, or other aspect of a proposed standard, factors beyond the technical merits of the proposal can influence community members' responses.

Because of this, rather than allowing a proposal's technical merits to speak for themselves, it is important to engage members of the development group to proactively build support. Attempt to identify allies who are willing to speak positively about the proposal rather than simply not-objecting to it. Talk to those who have expressed opposition and attempt to understand their reasons as it might be possible that their objections can be assuaged with minor changes to the proposal. Many standards bodies employ a voting system to decide issues, but it is the discussion ahead of that vote that determines the vote's outcome. A representative who can skillfully make the case for a particular proposal is arguably more valuable than having voting rights within the group – the latter only gives one vote, while the former can deliver a majority. In short, proposed solutions can benefit from the application of diplomacy as much as from the technical expertise that underpins those solutions.

Some caveats should be raised here. First, while most groups accept and encourage "hallway conversations", some groups want important discussions to occur via regular channels so that the process can be open and recorded. In these cases, one needs to be careful not to appear to be trying to undermine this process through too many backroom deals. Second, as with all diplomatic efforts, one must pick one's battles. A proposal which most parties feel to be flawed is unlikely to be accepted regardless of one's diplomatic skills, and pushing too hard for its acceptance can backfire. While all standards development groups have a political aspect to them, there almost always remains a strong pride in technological ability among participants, and groups can react negatively to any perceived attempt to overcome a technical gap via political maneuvering.

The ability to both provide solid technical proposals and apply diplomacy can be a powerful combination within a standards group. When an organization is selecting a representative to participate in a standardization effort on their behalf, the ideal candidate might not be the party with the greatest technical depth, but the one who can combine technical competency with soft skills such as diplomacy and negotiation. An impolitic presenter has doomed more than one otherwise technically acceptable proposal within standards bodies. This is not to say that everyone participating in a standardization effort is likely to be diplomatic - there are plentiful examples of the opposite. Social credit and reputation, as discussed earlier, can sometimes compensate for a lack of diplomatic skills when working within a group. However, especially when one is new to a standards group, soft skills can be an invaluable asset for a participant.

Recommendation: Actively build support for proposals within standards groups rather than relying on the technical merit of that proposal to speak for itself.

Recommendation: Organizations seeking to represent their interests within a standards body should select participants who have a mix of technical competence as well as diplomatic skills to improve the reception to proposals from that participant.

6.4 Active Participation

The two greatest factors in the perceived success of a standardization effort, as identified by Carl Cargill using the 1990 study by Martin Weiss and Marvin Sirbu, were an understanding of the standard's intended market (as discussed in section 5.4.2), and "the willingness of firms to commit written technological contributions to the standards committee". (C. F. Cargill 2011) The latter underscores a critical aspect of standardization efforts: standards groups, almost without exception, suffer from a dearth of available workers. At the same time, success of the group hinges on members writing, editing, and otherwise actively producing content and resources. Simply calling for others to work on a standard without demonstrating tangible effort of one's own is not only unhelpful, but can be seen as a lack of commitment by the party making that call. If one is invested in the outcome of a standardization effort, this investment needs to be supported by active participation in producing content for this effort.

There are many ways in which this participation can manifest. Obviously, writing sections of a draft specification is highly helpful, as is editing such work by others. Once aspects of a standard are codified, writing simple test implementations is another way the effort can be supported. Finally, many efforts involve some amount of research into related practices or technology, and someone needs to perform this research and report back to the group. While it should not be the job of any single party, even one strongly invested in the success of a standardization effort, to do all of this work, supporting the effort through active contributions not only moves the standardization process forward, but it demonstrates a commitment to the effort that is more likely to lead others to assist as well. In short, active participation breeds active participation.

Recommendation: Allocate sufficient resources to actively participate in the writing, research, editing, and/or prototyping work undertaken in the course of a standardization effort.

6.5 Schedules

The process of formal, open standards creation is often compared unfavorably to the speed and efficiency of “closed-shop” solution engineering. Indeed, any open standards creation effort is likely to take longer and be less well aligned with any one organization’s requirements than a small internal solution engineering effort. However, this comparison is misleading as the two activities have different goals. The goal of a standards development effort is not only to engineer a solution - it is to create a standard. A given engineering solution might address a particular technical problem or create a new capability, but only a standard offers uniformity of practice and enables bridging the gap between capability and information islands.

Solution engineering is only one part of the standard development process. Standardization efforts require building consensus around a solution among multiple parties, many of whom approach the effort from different perspectives. Once consensus around a solution has been achieved, group members must craft a document that can be universally understood and produce identical behavior in highly complicated technical systems. This document then usually needs to be submitted to a (sometimes lengthy) formal vetting process. All of these steps have an impact on both the timeliness and agility of standards bodies, but are necessary for the production of a usable technical standard. While solution engineering may be the activity that interests some participants most[14], it is only through all of the described steps that a proposed solution is turned into a viable technical standard.

The steps described above are one reason that standardization efforts can have long timelines. However, there are other factors that can further delay these efforts. Most participants in a standardization effort have competing calls for their time, which means timelines can slip as milestone are missed. Consensus building is often a source of uncertainty in producing time estimates and it can be difficult to predict which issues are likely to be contentious. Moreover, the discovery of a problem in a standard draft during vetting can force the standard to repeat review steps, adding months to the timeline. As a result, the time needed for the creation of a standard can vary significantly, often for reasons that are difficult to anticipate, and almost always is longer than scheduled.

Participants need to be aware of the probability that timelines will slide, sometimes by many months. There is usually little any individual participant can do about this slippage. Parties that

[14] The author is aware of a few instances where a standardization effort was convened by someone only interested in the development of an engineering solution and who saw a standardization effort as a way to cheaply “crowd-source” the solution design. Invariably, the convener rapidly became frustrated by the pace of work and repeated divergence from their own priorities.

are dedicated to the successful completion of a standardization effort should be prepared for deadlines to slip and allocate sufficient resources to remain engaged if/when this happens.

Recommendation: Expect that any standard creation effort will take longer than scheduled and allocate an appropriate level of resources.

6.6 In-Person Meetings

One aspect of active participation that deserves attention is the importance of physically attending in-person meetings. Standards bodies usually host a small number of in-person meetings each year. For international standards bodies, these meetings are usually physically dispersed around the world to avoid regularly inconveniencing participants from certain regions.[15] This usually means that attending a group's meetings can require periodic international travel, accompanied by commensurate costs.

When actively engaged in a standardization effort, the ability to attend in-person meetings is important. In addition to enabling direct participation in scheduled discussions, attending these meetings allows participants to engage in hallway conversations, which are often invaluable in developing technical solutions, creating strategies for progress, and for diplomatic support building. Virtually all groups attempt to accommodate participants who cannot physically attend meetings through remote meeting technology, but today such technology is still not equal to physical presence. Remote participants cannot participate in discussion to the same degree as their present colleagues and generally are unable to participate in hallway conversations at all. While the cost of international travel can produce "sticker-shock", that price needs to be placed in the context of the overall investment in supporting the effort and seeing it to a successful conclusion. Compared to one's total investment in a standard creation effort, and noting the generally high productivity and value achieved during in-person meetings, physical attendance in these meetings is not only justifiable, but can provide very high return on investment.

Of course, not all levels of participation justify the cost of physical presence at these meetings. As noted earlier, it is beneficial to build reputation and social credit in standards groups through ongoing involvement, even when there is no activity in which the participant or their sponsor are specifically invested. During such interludes, remote participation in in-person meetings can be sufficient to maintain engagement while minimizing costs. There is usually at least one meeting that is relatively close to any given participant and physical attendance of a local meeting can further demonstrate commitment and engagement to the body's activities, which can build reputation. However, when supporting a priority effort (rather than just maintaining low-level involvement), this support requires some level of investment in travel to participate in person.

Recommendation: When actively participating in the development of a particular standard, allocate sufficient resources to allow physical attendance at in-person meetings of the standards body.

[15] With regard to the often-voiced assertion that these meetings are primarily an excuse for members to visit exotic locations, it should be noted that all exotic locations look largely the same from the inside of hotel conference rooms where participants end up spending most of their time. Standards bodies generally run very full schedules at these meetings to maximize productivity.

7 Conclusion

Technical standards can play an important role in defending today's complex, distributed IT infrastructures from attack. The use of standards limit attackers' options, while allowing defenders to concentrate their attention and resources on a smaller number of attack vectors and to develop more cooperative and sophisticated defensive solutions. New technologies and practices are constantly arising, forcing vendors to respond with ad-hoc mechanisms. This hampers not only those who must defend enterprises that include such devices, but hampers those who develop products and technologies to assist in that defense. As a result, there is an ongoing need to identify the cases where non-standard behavior leads to security challenges and seek to address them through the development of new technical standards.

The US government has an important role to play in addressing this gap. The US government is a uniquely capable participant in standards development efforts due to its market leverage, its broad perspective, its deep expertise in cybersecurity, and its responsibility in several sectors where security, safety, and reliability are critical issues. Engagement by the US government in the creation of standards is not only beneficial to the standards development process, but can contribute significantly to the creation of a more secure, safe, and reliable national cyber infrastructure.

Of course, it is not practical for the US government to engage in every standardization effort or have participants in every standards body and industry consortia. Instead, government agencies need to be selective in their engagements to make optimal use of their resources. Towards this end, it can be beneficial to consider existing priorities and use these to guide participation in standardization effort, be they joining existing efforts or proposing new ones. In the case of the latter, care should be taken in identifying a body or consortia that aligns with the proposer's understanding of the problem that needs to be addressed, and which can help bring a usefully broad and engaged set of participants to the table. If a particular body or consortia is engaged with some frequency, this likely indicates a good venue in which to continue low-level participation even between specific efforts of interest so as to build social credit in that group. In short, representatives of the US government should consider their own priorities and then identify situations where the development of a standard is an appropriate and beneficial tool to meet that priority. This paper cannot speak to US government priorities, but does provide recommendations that can help inform where standardization can be practical and beneficial, and then help that standardization effort to successfully meet US government needs.

The creation of standards is similar to many other engineering efforts, but also has some unique requirements and challenges. Moreover, working in open standards creation bodies requires different skills and practices than those required by a private engineering effort. This paper looks at some of the characteristics that are of particular significance in standard creation and, through the consolidation of input from multiple, highly experienced participants in standards communities, provides recommendations that can help reduce the risk of failure in such efforts. Many factors in standards creation remain beyond the control of any one participant, and thus no checklist is capable of guaranteeing the success of such endeavors. However, through the guidance captured in this paper, US government efforts to create a more secure cyberspace through the creation of technical standards will have a greater chance of success.

List of References

Cargill, Carl F. 2011. "Why Standardization Efforts Fail." *Standards* 14 (1).

Cargill, Carl. 2011. "On standards as change agents - From a pragmatic business perspective ." *Standardization and Innovation in Information Technology (SIIT).* Berlin.

Chopra, Aneesh, Miriam Sapiro, and Cass R. Sunstein. 2012. *Principles for Federal Engagement in Standards Activities to Address National Priorities (OMB Memo M-12-08).* Washington D.C.: Executive Office of the President.

Defense Advanced Research Projects Agency. n.d. *Paving the Way to the Modern Internet.* Defense Advanced Research Projects Agency. Accessed February 9, 2017. http://www.darpa.mil/about-us/timeline/modern-internet.

Folk, Chris, Dan C. Hurley, Wesley K. Kaplow, and James F. X. Payne. 2015. *The Security Implications of the Internet of Things.* Fairfax, VA: AFCEA International Cyber Committee.

Hogan, Michael, and Elaine Newton. 2015. *Report on Strategic U.S. Government Engagement in International Standardization to Achieve U.S. Objectives for Cybersecurity (NISTIR 8074).* Gaithersburg, MD: National Institute of Standards and Technology.

Mann, David, Stuart S. Shapiro, and Deb Bodea. 2014. "Bilateral Analysis of Information Sharing Efforts: Determining the Expected Effectiveness of Information Sharing Efforts." *Proceedings of the 2014 ACM Workshop on Information Sharing & Collaborative Security* . New York, NY.

US Department of Homeland Security. 2009. "Chemical Facility Anti-Terrorism Standards Risk-Based Performance Standards Guidance." May. Accessed December 16, 2015. Chemical Facility Anti-Terrorism Standards Risk-Based Performance Standards Guidance.

—. 2015. *Science and technology | Homeland Security.* US Department of Homeland Security. Accessed November 10, 2015. http://www.dhs.gov/science-and-technology.

US National Institute of Standards and Technology. 2015. *Computer Security Division Homepage.* US National Institute of Standards and Technology. September 14. Accessed November 10, 2015. http://www.nist.gov/itl/csd/index.cfm.

US National Security Agency/Central Security Service. 2015. *Research - NSA/CSS.* US National Security Agency/Central Security Service. February 09. Accessed November 10, 2015. https://www.nsa.gov/research/.

Weiss, Martin BH, and Marvin Sirbu. 1990. "Technological choice in voluntary standards committees: An empirical analysis." *Economics of Innovation and New Technology* 1 (1-2): 111-133.

Wi-Fi Alliance. n.d. *Wi-Fi Alliance.* Wi-Fi Alliance. Accessed November 19, 2015. http://www.wi-fi.org/.