"Early Stage"
Coordinated Vulnerability Disclosure Template
Version 1.1[1]

NTIA Safety Working Group
December 15, 2016

Executive Summary

Collaboration between technology providers and security researchers has become an important part of good information security. As security researchers increasingly discover vulnerabilities in organizations' technology, those organizations benefit from having a process in place for working with the researcher to understand and mitigate the risk. To help foster this collaboration across the digital ecosystem, the National Telecommunications and Information Administration (NTIA) convened a multistakeholder process to address principles and practices around security researcher disclosure.[2]

This document reflects the work of the "Safety" working group, which focused on the initial steps an organization can take to improve collaboration. It was developed by experts in an open, transparent fashion, with diverse participation from industry, government, and the security community. Much of the discussion targeted the safety-critical industry, in which the potential for harm directly impacts public safety or causes physical damage (e.g., automobiles or medical devices), but the lessons are easily adaptable by any organization that builds or maintains its own software or systems.

In this report, we discuss why security disclosure is important, particularly for safety-critical industries that are becoming more and more dependent on software and digital systems. We present a template disclosure policy, explain the different sections, and offer a sample policy for "Acme Corp." At the end of this document, we walk through critical issues any organization should consider when developing a security disclosure policy of its own.

[1] The Working Group is soliciting public comment on this draft, and intends to issue a revised version in early 2017. Please send feedback to: afriedman@ntia.doc.gov to pass along to the working group. *The deadline for feedback is February 15, 2017.*

[2] More information on NTIA's open Multistakeholder Process to Promote Collaboration on Security Research Disclosure that process is available at https://www.ntia.doc.gov/other-publication/2016/multistakeholder-process-cybersecurity-vulnerabilities.

Introduction: Disclosure and Safety

Safety-critical systems are increasingly dependent on software, and therefore increasingly subject to software security issues. Coordinated vulnerability disclosure directs energy and attention into improving the safety and security of systems and software for the overall population. Compared with traditional IT systems, manufacturers of safety-critical systems have a higher consequence of failure and relatively less experience with vulnerability disclosure. High trust, high collaboration interactions come from understanding mutual expectations and perspectives.

We define "safety-critical industries" as those in which the potential for harm directly impacts humans – for example, an automobile, an embedded medical device (such as a pacemaker or insulin pump), or carbon monoxide detectors. Compared with Enterprise IT, Safety has several differences that must be appreciated and accounted for. These differences[3] will impact more than just disclosure policies and actions (by multiple stakeholders); manufacturers should also consider how design choices will limit or grant capabilities to react to newly discovered vulnerabilities.

- Consequences: When software is a dependency for safety-critical systems, consequences of security failure may manifest in direct, individual harm, including loss of life. Impacts from wide-scale harm can shatter confidence in the firm or the market, and can damage trust in government and its role safeguarding citizens through oversight and regulation.

- Adversaries: Different adversaries have different goals, motivations, methods, and capabilities. While some adversaries may be deterred by potential harm from safety-impacting systems, others may seek these systems out. For instance, ideological actors may wish to inflict harm, and criminal groups may suspect owners will pay higher ransoms.

- Composition: Some components in Internet of Things devices, including safety systems, are not found in typical IT environments. Elements such as sensors, programmable logic controllers, low power chips, embedded controllers, limited battery life, etc., limit capabilities available to the manufacturer in design and response.

- Economics: Components for safety systems may require a high degree of resourcing to protect and have a very low cost of goods, and profit margins may also be smaller. Security capabilities for million-dollar data centers are likely cost prohibitive in 42 cent microchips, for example.

- Context and Environment: Safety-critical systems often exist in unique operational, environmental, physical, network, immediacy/real-time, and legal contexts. For

[3] I Am the Cavalry. "6 Differences in Internet of Things and Cyber Safety." Available at: https://www.iamthecavalry.org/iotdifferences/

instance, a pacemaker is implanted in a human body, has no IT staff, must respond immediately, has no bolt-on security measures, and carries strict regulatory requirements.

- Timescales: Timescales for design, development, implementation, operation, and retirement are often measured in decades. Response time may also be extended because of composition, context, and environment. Safety systems in design today may be with us for 10, 20, 40, or more years.

Vulnerability disclosure and remediation in cyber safety contexts should be handled with both due haste and due care. Researchers may be more reluctant to disclose if they know a vulnerability has not been (or cannot be) fixed. On the other hand, the prospect of high consequence failures may motivate action. Remediation urgency can preserve safety, life, and trust; at the same time, validation and verification avoid unintended consequences, which can increase risk. Decisions considered insecure for a web application may be appropriate for an implanted medical device. Any hard deadline for disclosure or remediation may both be too long and too short to safely address security vulnerabilities in safety-critical systems.

We believe Coordinated Vulnerability Disclosure is especially important – and urgent – for safety-critical industries. DMCA research exemptions[4], which remove significant legal barriers to security research on cars and medical devices, went into effect in late October 2016. With softened fear of legal concerns, higher numbers of researchers are likely to engage in vulnerability research and disclosure in safety-critical industries. Organizations in those sectors should understand how the security research community may want to engage and equip themselves with a flexible set of tools to successfully collaborate and improve security.

Disclosure Policy: The First Steps

Stakeholders representing a range of interests in this community recommend a considered approach that starts small to build experience, confidence, trust, and capacity. Firms contemplating their first steps into Coordinated Vulnerability Disclosure have many resources and references from multiple sources available to consult as they develop their programs. This journey has taken many years for even the most sophisticated technical organizations.

What follows is a simple framing of what an "early stage" coordinated disclosure program might look like. Below, we present a template of what a successful, lightweight, and adaptable disclosure policy might look like and then highlight some notable issues in developing such a policy. We also present a sample disclosure policy.

[4] US Copyright Office. "Exemption to Prohibition on Circumvention of Copyright Protection Systems for Access Control Technologies" 80 FR 65944 (2015). Available at: https://www.federalregister.gov/documents/2015/10/28/2015-27212/exemption-to-prohibition-on-circumvention-of-copyright-protection-systems-for-access-control

There are many resources on how to think about vulnerability disclosure and handling, including ISO/IEC Standards 29147 and 30111.[5] For more information, two other documents produced by stakeholders in the NTIA process may be of further interest: "Vulnerability Disclosure Attitudes and Actions: A Research Report,"[6] with more background on security disclosure, and "Guidelines and Practices for Multi-party Vulnerability Coordination,"[7] for organizations facing more complex disclosure challenges.

[5] ISO/IEC 29147 "Vulnerability Disclosure" (2014) http://www.iso.org/iso/catalogue_detail.htm?csnumber=45170. This standard is publicly available at: http://standards.iso.org/ittf/PubliclyAvailableStandards/index.html. ISO/IEC 30111 "Vulnerability Handling Processes" (2013) can be be found at http://www.iso.org/iso/catalogue_detail.htm?csnumber=53231

[6] "Vulnerability Disclosure Attitudes and Actions: A Research Report" (2016). This stakeholder-drafted report is available at https://www.ntia.doc.gov/files/ntia/publications/2016_NTIA_A_A_vulnerability_insights_report.pdf

[7] "Guidelines and Practices for Multi-party Vulnerability Coordination" (2016). This stakeholder-drafted report is available at https://www.first.org/global/sigs/vulnerability-coordination/multiparty

Template Disclosure Policy

The first step an organization should take is to develop a Coordinated Vulnerability Disclosure policy. We urge the creation/use of a simple, short document. These can fit on a *single, readable* page. Many organizations, including automakers and medical device makers, have already done this, leveraging the template below.

Brand Promise

Objective: To demonstrate a clear, good faith commitment to customers and other stakeholders potentially impacted by security vulnerabilities.
Audience: Customers and the market
Tone: Committed, concerned, and open. For instance, "The safety & security of our customers is important to us…"
Content: Assure customers and the market that safety and security is important. Describe what work has already been done as well as future commitments. A message to the vulnerability reporter can serve as outreach and can build trust up front. For instance, "we are undertaking this program to give security researchers a point of contact so they can directly submit their research findings, which can then be remediated in a prioritized and efficient way."

Initial Program and Scope

Objective: To outline which systems and capabilities are "fair game" versus "off limits" for the initial program, which will evolve as capacity and confidence change.
Audience: Vulnerability finders and reporters
Tone: Set a reasonable initial phase to build capacity.
Content: Declaration of explicit and/or implicit scope, and optionally define what's out-of-scope. Explicit scope sets an expectation for what vulnerabilities will be addressed from reports, such as models/years and versions as well as duration. Implicit scope, such as recognition and/or reward, allows a degree of throttling of program participation (see below), and can be expanded over time as well. Optionally, defining what is "out-of-scope" can prevent unintended harm from good faith research, though a long list may dissuade research.

"We Will Not Take Legal Action If…"

Objective: To assure that vulnerability finders and reporters of good faith receive responses to their good faith acts.
Audience: Vulnerability finders and reporters
Tone: Non-threatening, inviting, and reasonable, using language accessible to individuals without a legal background or representation. Affirmative language tends to be better received than prohibitive, with some key exceptions such as "testing implanted devices is excluded."
Content: Clear, unambiguous statements that guide researchers' good faith efforts. This section should tell researchers what activities will and won't result in legal action, in a way that is

evergreen and is very unlikely to change. This section can also outline safety consequences from deviating.
Other Considerations: This section should contain legal postures, rather than preferences and priorities, which will come later. Parties should account for applicable state, local, and national/federal laws.

Communication Mechanisms and Process

Objective: To clearly identify communication mechanisms and reasonable acknowledgement timeframe.
Audience: Vulnerability finders and reporters
Tone: Reasonable for the initial information exchange
Content: Define a mechanism for submission and reporting, including security precautions (such as a PGP encryption key) and requirements for completeness of submission (different from a legal posture). Many organizations prefer a secure web form. This section should also set expectations for when the researcher can expect to receive acknowledgement of the submission and how future engagement/communication will take place. This section can outline conflict resolution mechanisms and roles and responsibilities.

Nonbinding Submission Preferences and Prioritizations

Objective: To set expectations based on priorities and submission volume, rather than based on legal objection or restriction.
Audience: Vulnerability finders and reporters
Tone: How bugs will be triaged/prioritized
Content: This section is a living document that sets expectations for preferences and priorities, typically maintained by the support and engineering team. This can outline classes of vulnerabilities, reporting style (crash dumps, CVSS scoring, proof-of-concept, etc.), tools, etc. Too many preferences can set the wrong tone or make reporting findings difficult to navigate. This section also sets expectations to the researcher community for what types of issues are considered important or not.

Versioning

Objective: To track the evolution of the policy.
Audience: Vulnerability finders and reporters
Tone: Organized to help the researcher understand potential future changes, as well as past adjustments to the policy.
Content: This optional section can help the reader understand how the policy has evolved, and how it might evolve in the future. See “Changing the Disclosure Policy” below.

Sample Vulnerability Disclosure Policy Template

ACME Corp.

Brand Promise

ACME Corp., the leading manufacturer of embedded software widgets, is committed to ensuring the safety and security of our customers. Toward this end, ACME is now formalizing our policy for accepting vulnerability reports in our products. We hope to foster an open partnership with the security community, and we recognize that the work the community does is important in continuing to ensure safety and security for all of our customers.

We have developed this policy to both reflect our corporate values and to uphold our legal responsibility to good-faith security researchers that are providing us with their expertise.

Initial Program and Scope

Initial Scope

ACME's Vulnerability Disclosure Program initially covers the following products:

- ACME Widgetsoft 3.1
- ACME Widget Module A
- ACME Widget Module B
- ACME Widget Controller
- ACME Widget Ethernet Gateway Module

While ACME develops a number of other products, we ask that all security researchers submit vulnerability reports only for the stated product list. We intend to increase our scope as we build capacity and experience with this process.

Researchers who submit a vulnerability report to us will be given full credit on our website once the submission has been accepted and validated by our product security team.

We Will Not Take Legal Action If…

Legal Posture

ACME Corp will not engage in legal action against individuals who submit vulnerability reports through our Vulnerability Reporting Form. We openly accept reports for the currently listed ACME products. We agree not to pursue legal action against individuals who:

- Engage in testing of systems/research without harming ACME or its customers.
- Engage in vulnerability testing within the scope of our vulnerability disclosure program and avoid testing against [ex. website].
- Test on products without affecting customers, or receive permission/consent from customers before engaging in vulnerability testing against their devices/software, etc.
- Adhere to the laws of their location and the location of ACME. For example, violating laws that would only result in a claim by ACME (and not a criminal claim) may be

acceptable as ACME is authorizing the activity (reverse engineering or circumventing protective measures) to improve its system.
- Refrain from disclosing vulnerability details to the public before a mutually agreed-upon timeframe expires.

Communication Mechanisms and Process

How to Submit a Vulnerability

To submit a vulnerability report to ACME's Product Security Team, please utilize the following form <link to vulnerability reporting form>[8].

Nonbinding Submission Preferences and Prioritizations

Preference, Prioritization, and Acceptance Criteria

We will use the following criteria to prioritize and triage submissions.

What we would like to see from you:

- Well-written reports in English will have a higher chance of resolution.
- Reports that include proof-of-concept code equip us to better triage.
- Reports that include only crash dumps or other automated tool output may receive lower priority.
- Reports that include products not on the initial scope list may receive lower priority.
- Please include how you found the bug, the impact, and any potential remediation.
- Please include any plans or intentions for public disclosure.

What you can expect from us:

- A timely response to your email (within 2 business days).
- After triage, we will send an expected timeline, and commit to being as transparent as possible about the remediation timeline as well as on issues or challenges that may extend it.
- An open dialog to discuss issues.
- Notification when the vulnerability analysis has completed each stage of our review.
- Credit after the vulnerability has been validated and fixed.

If we are unable to resolve communication issues or other problems, ACME may bring in a neutral third party (such as CERT/CC ,ICS-CERT, or the relevant regulator) to assist in determining how best to handle the vulnerability.

Versioning

This document Version 1.1 was created 15-December-2016. [We update or renew this policy every 90 days.] Any updates will be noted below in the version notes.

[8] For an example of a secure web form, see cert.org's Vulnerability Reporting Form: https://vulcoord.cert.org/VulReport/form

Issues to Consider in Writing a Disclosure Policy

Defining Vulnerability Disclosure Program Scope

Any newly implemented vulnerability disclosure program may need to deal with a large, unanticipated volume of submissions. In the early stage, report volume can be reduced through explicit or implicit scoping in the disclosure policy. This has the effect of focusing researchers on the specific type of disclosure items the company is prepared to respond to while it builds capacity and experience.

For example, submissions could be explicitly scoped by limiting the program to the following:

- Only specified product model years
- Only select product make/model/year
- Only particular types of vulnerabilities

Implicit scoping may be influenced by the type, structure, and scale of incentives that may be awarded to researchers, if any incentives are used at all. Asking researchers to focus in a particular area for finding security issues is one way of scoping. Another limit to a program's scope may come from the reward structure. A Coordinated Vulnerability Disclosure Program with no reward program is likely to attract altruistic individuals or hobbyists who want to share their findings with the company, but are not looking for a reward. Adding modest recognition and/or a reward to the program could expand the scope to increased researcher participation. Rewards such as providing recognition on a wall of fame board or awarding a challenge coin and/or branded merchandise attracts some researchers to find vulnerabilities. Larger financial rewards will attract researchers as well, and will be less likely than the previous scenarios to limit the response from the research community.

Researchers are motivated to understand security flaws for a wide range of reasons, from a desire to solve an interesting problem to a desire to protect other users; the table below illustrates some of the diverse types of motivations researchers may have. Benefits of narrowing the scope and/or having no financial incentive for reporting vulnerabilities include: limiting the number of reported vulnerabilities; and attracting researchers who may have more patience and/or less motivation to disclose during conference presentations (with submission deadlines), the dates of which could conflict with the organizational process.

Table 1 - Diverse Motivations of Security Researchers[9]

Researchers Motivations	***Description***
Protect	Wants to make the world a safer place. May be more sensitive to realities affecting safety.
Puzzle	Tinkerers, curiosity, hobbyists. Driven by 'How does this work?'
Prestige/Pride	Recognition, making a name, conference & media visibility.
Profit/Professional	Seeking monetary reward and/or making a living off it.
Politics/Patriotism/Protest	Ideological or principled. E.g. Civil liberties. Strongly pro- or anti- causes or organizations.

In summary, an organization can use explicit and implicit scoping mechanisms to match its capacity to implement its disclosure program. As the organization builds capacity and experience through responses to vulnerability disclosures, it can scale its program accordingly. With maturation of the organizational response capabilities, explicit and implicit scope limitations may be relaxed so that more useful disclosures might be obtained. Additionally, vulnerabilities that fall outside the program scope may still deserve appropriate consideration and response. Programs should be prepared for such a contingency, and avoid turning away well-intentioned finders who are aware of a vulnerability in a system, even if it's out of scope of the current policy.

Changing the Disclosure Policy

As with any policy, at some point, it may need to be changed or modified. The side effect of changing the disclosure policy is that it can make things difficult for researchers to navigate and difficult for vendors to track, or can cause researchers to lose confidence in vendor promises. As such, we recommend minimizing changes if possible. While that may not be possible, the legal protections offered to researchers should not change much over time if trust is to be maintained.

Given that policies may change, some strategies to maintain trust include:

- Be transparent – explain why the disclosure policy is changing
    - Accept feedback on changes / listen to the community
- Explicit duration of any given policy: This policy is good until <pre-defined date>.
- Include version control
    - For any change made; archive prior versions (consider archiving ALL versions on the organization's site)

[9] I Am The Cavalry. "5 Motivations of Security Researchers." Available at: https://www.iamthecavalry.org/motivations/

    - Avoid abrupt or erratic changes in the policy, and provide updates on consistent time periods
- Consider allowing researchers to enroll, and become grandfathered into a given policy version
    - This puts a lot of responsibility on to the researcher and the vendor to track which policy version is being used
    - [Light version: have a feed or email list for updates]
- Include explicit caveats about how the policy will change
    - This may result in a very long and complicated policy
    - Black lists will invariably grow
    - Potential solutions: white listing (allowed) over black listing (disallowed)
- Declare certain parts of the policy immutable, *particularly legal protections and promises*
    - Have a baseline – everything above this point is the basics, won't change
        - Baseline = white list (allowed)
        - Consider tying in with brand promise
        - Should reflect high level goals of program or impacts of vulnerabilities rather than technical approaches
        - Changes to white list (adding or removing) should be rare and accompanied with an explanation for the change
    - Here is the section that we may change – establish what might trigger a change
        - Changing = black list
        - May be used to throttle common or "low effort" vuln reports
        - May change as a result of enhanced security engineering / QA process
        - May be used to shift the focus to the newest or riskiest product
    - Can encourage researchers to check back, and archive which version they started the research against (in good faith) to grandfather themselves in
    - Can subscribe to an RSS feed of updates

Resolving these issues will help inspire confidence among researchers, helping to promote the success of the policy.

Restrictions on Disclosure

Researchers do not create vulnerabilities. The fact that one researcher does not disclose its existence does not guarantee that another will not find it – or has not already found it. Finders may have reasons to want to disclose the vulnerability publicly, including the range of motivations discussed above. A managed disclosure situation is preferable to one without control. Vendors may want to express preferences on when finders publicly talk about vulnerabilities. A few options are:

Do not publicly disclose:

1. Until it is fixed

2. Until a particular timeframe after first submission
3. Until after giving the organization X days of notice
4. Mutually agreed-upon (or negotiated) timeline (as discussed above, different technologies or sectors may have different timelines), which may be adjusted as part of the process with the disclosing party. (Note: Communication with the disclosing researcher is critical in this part of the process because communication is the only way a researcher will know progress is occurring and the organization is taking the issue seriously)

There are strong pros and cons for denying researchers any ability to go public. For example, if an organization states that "no disclosures can happen until the bug is fixed," then there may be less risk of exploitation, but there may also be risk that some researchers will not participate. What if they fear a vendor "sitting" on a bug?

- What if the fix takes 5 years?
- Some researchers may expect very fast turnarounds for bugs, but some affected industries can't turn on a dime.

Because reasonable people can disagree on the method and manner of public disclosure, it may also be prudent to have a defined path of escalation and mediation and to factor in the appropriate guidance/participation from the regulator of jurisdiction or relevant parts of governments (e.g. US FDA or NHTSA – and US DHS-ICS-CERT). For something like an embedded medical device, the FDA may be best poised to determine impact to patient care across the ecosystem – as well as the optimal safety communication strategy.