5.7 Disclosure Timing

When you say it's gonna happen now,
When exactly do you mean?
See I've already waited too long
And all my hope is gone

-The Smiths, "How Soon is Now?"

How long is "long enough" to respond to a vulnerability? Is 45 days long enough? Is 90 days too short? Is 217 days unreasonable? Three years? Talk among yourselves. We can wait.

As with so many questions that arise in the CVD process, there is no single right answer. So rather than trying to solve an underspecified set of inequalities, let's have a look at some of the factors that tend to play into timing choices. This will give us an opportunity to see where some of the variability comes from.

Conference Schedules and Disclosure Timing

Conference schedules often drive researcher timelines. This is a big one. There is a rhythmic cycle to the vulnerability disclosure calendar. Black Hat [1] and DEF CON [2] happen in early August every year. Usenix Security [3] is usually right after that. The RSA Conference [4] is in the late winter or early spring. CanSecWest [5] is in the spring. Smaller conferences are scattered in between. Many of these conferences rely on presenters describing novel attack methods in varying degrees of detail. However, in order for researchers to analyze, develop, and demonstrate those techniques, vulnerabilities are often uncovered in extant products. That means that coordinating the disclosure of the vulnerabilities they've found is a common part of the conference preparation process for presenters. The CERT/CC often observes an increased rate of vulnerability reports a few months in advance of these conferences. Vendors would do well to be aware of these schedules and be prepared to respond quickly and appropriately to seemingly inflexible deadlines for disclosure.

Vendor Reputation and Willingness to Cooperate

Vendors that are perceived to treat vulnerability reporters poorly or that are perceived to be slow or unresponsive may find themselves being left to discover reports of vulnerabilities in their products at the same time as the public becomes aware of them. CVD is a social process, remember? And the game is played over and over, by players who share knowledge between rounds.

Declarative Disclosure Policies Reduce Uncertainty

Avoiding surprise was one of the principles in Section 2. To that end, explicitly declared policies (from both researchers and vendors) are a good thing. Expected disclosure timing is an important question to ask whenever a report is received. Sometimes the reporter or coordinator acting on the reporter's behalf has a standing policy of X days with no exceptions. Other reporters may be more flexible. If in doubt, ask.

All Disclosure Agreements Are Contingent

When vendors, reporters, and/or coordinators negotiate and agree to a release timeline for a vulnerability, they may behave as if they've reached some state of détente with the world. But that's an illusion. True, they may have reached an agreement with each other, but it ignores another relevant role in the disclosure process: the adversary. Adversaries do not care whether the vendor plans to release the patch in a month, whether they need more time to test it prior to release, whether the reporter wants to break the news at a conference event, whether the reporter hopes to get paid for the work, or any other reason that reporters and vendors may have for agreeing to the terms they came to. Furthermore, because the vulnerability's existence is an observable fact in the world, anyone else who happens to notice it might also choose to disclose that knowledge on their own terms without being party to any existing embargo agreements. Therefore it's important for vendors, reporters, and coordinators alike to recognize that all disclosure embargo agreements are necessarily contingent on circumstances beyond the control of the parties involved; and that those circumstances are not simply random events but may be controlled by actors who are indifferent to their concerns.

Diverting from the Plan

Je n'ai jamais eu un plan d'opérations.

-Napoleon Bonaparte

Plans are one thing, but reality sometimes disagrees with our assessment of it. Breaking a previous disclosure timeline agreement is sometimes necessary when events warrant. Below we cover a few reasons to release earlier or later than planned.

Reasons to release early include:

- Evidence of active exploitation
- Vendor fails to respond, is not acting in good faith, or denies the existence of a vulnerability
- Vulnerability is known to be discovered by adversaries, so the race to defend vulnerable systems is more focused
- All known users have been notified and patched (usually via private channels)

Reasons to hold back release include:

- Vendor not ready with fix, but continuing to make progress and is acting in good faith
- Vulnerabilities with severe impact, especially those affecting safety-critical or critical infrastructure
- Cases where new information is found late in the process, for example that there are important but previously unrecognized dependencies that alter the impact of the vulnerability or patch deployability

In cases that divert from the planned disclosure date, it sometimes helps to seek the opinion of a neutral third party for advice on how to proceed. Finders, reporters, and vendors can each have valid yet conflicting perspectives on what the best course of action might be. Coordinator organizations are often able to help resolve conflicts by taking a neutral approach to the situation and advising one or more parties in light of their prior experience.

Releasing Partial Information Can Help Adversaries

When considering what information to release about a vulnerability, our advice is "Don't tease." Our experience shows that the mere knowledge of a vulnerability's existence in a feature of some product is sufficient for a skillful person to discover it for themselves. Rumor of a vulnerability draws attention from knowledgeable people with vulnerability finding skills—and there's no guarantee that all those people will have users' best interests in mind. Thus, teasing the existence of a vulnerability in a product can sometimes provide an adversarial advantage that increases risk to end users.

References

1. Black Hat, "Black Hat," [Online]. Available: https://www.blackhat.com/. [Accessed 23 May 2017].
2. DEF CON, "DEF CON," [Online]. Available: https://www.defcon.org/. [Accessed 23 May 2017].
3. USENIX, "USENIX Security Conferences," [Online]. Available: https://www.usenix.org/conferences/byname/108. [Accessed 23 May 2017].
4. RSA, "RSA Conference," [Online]. Available: https://www.rsaconference.com/. [Accessed 23 May 2017].
5. CanSecWest, "CanSecWest Vancouver 2018," [Online]. Available: https://cansecwest.com/. [Accessed 23 May 2017].