The U.S. Vulnerabilities Equities Process

An Economic Perspective

Tristan Caulfield Christos Ioanndis David Pym

October 24, 2017

University College London

Aston Business School

Alan Turing Institute

Introduction

Vulnerabilities

Vulnerabilities in software are security risks — they can cause harm to individuals, businesses, and governments.

But they're also useful to government for law enforcement, intelligence, and national security purposes. Apple v FBI, 'playpen', EU sharing.

Trade-off between security and government requirements.

Increasing demand for vulnerabilities

Prices for vulnerabilities have been steadily increasing:

2007	Difficult to sell exploits, not that lucrative
2012	30-60k USD for Android, 100-250k for iOS
Today	1.5 million USD for iOS

These are black/grey hat prices. Some of these vulnerabilities probably end up with governments.

The Vulnerabilities Equities Process

The VEP is a process that the U.S. government uses to decide to retain or disclose vulnerabilities it becomes aware of.

What do we know? Very little.

The paper aims to:

- help understand how different factors influence the decision
- give insight into how good decisions might be made
- provide a way for government decisions to be evaluated

It's not about being normative.

VEP Background

A bit of history

2008: President Bush signs a directive creating a Comprehensive National Cybersecurity Initiative. Required development of a plan for coordinating the 'application of offensive capabilities to defend US information systems'

2010: This CNCI led to the production of the VEP document.

2014: The EFF acquired a redacted version of this document.

2014: In response to Heartbleed, Michael Daniel, White House Cybersecurity Coordinator, writes a blog post about the factors used in the decision-making process. But — 'there are no hard and fast rules'.

From the VEP document

This document establishes policy and responsibilities for disseminating information about vulnerabilities discovered by the United States Government (USG)...

From the VEP document

This document establishes policy and responsibilities for disseminating information about vulnerabilities discovered by the United States Government (USG)... This policy defines a process to ensure that dissemination decisions regarding the existence of a vulnerability are made quickly, in full consultation with all concerned USG organizations, and in the best interest of USG missions of cybersecurity, information assurance, intelligence, counterintelligence, law enforcement, military operations, and critical infrastructure protection.

Recommendations

However, we don't know much about the decision-making process itself.

Schwartz and Knake, 2016: 'the principles guiding these decisions, as well as a high-level map of the process that will be used to make such decisions, can and should be public'

EFF, 2016: 'We think the government should be far more transparent about its vulnerabilities policy. A start would be releasing a current version of the VEP without redacting the decision-making process...'

Factors

Factors

These are the factors from the Daniel blog post:

- How much is the vulnerable system used in the core internet infrastructure, in other critical infrastructure systems, in the U.S. economy, and/or in national security systems?
- Does the vulnerability, if left unpatched, impose significant risk?
- How much harm could an adversary nation or criminal group do with knowledge of this vulnerability?
- How likely is it that we would know if someone else was exploiting it?

Factors

- How badly do we need the intelligence we think we can get from exploiting the vulnerability?
- Are there other ways we can get it?
- Could we utilize the vulnerability for a short period of time before we disclose it?
- How likely is it that someone else will discover the vulnerability?
- Can the vulnerability be patched or otherwise mitigated?

Not much else

No information about how these factors are quantified (or if that is even attempted).

No information about how the factors are combined, or how they relate to each other.

Lack of ‘hard and fast rules’ — possibly ad-hoc, case-by-case decisions.

Timing Problem

Typically, the discussion is about whether to disclose or not. We think it's more useful to think about *when* to disclose. We can think of this as a timing problem:

$$\max_T V_T = B_T - C_T,$$

We want to find the best time T to disclose, which maximises the value to the government.

Let's consider how the factors will affect the timing.

Extent of use

How much is the vulnerable system used in the core internet infrastructure, in other critical infrastructure systems, in the U.S. economy, and/or in national security systems?

Effect. Greater use means potentially greater harm, which would accelerate disclosure, but it is also potentially more useful to the government, which would delay disclosure.

Risks and harm

Does the vulnerability, if left unpatched, impose significant risk? How much harm could an adversary nation or criminal group do with knowledge of this vulnerability?

Effect. Aversion to substantial harm from single events will accelerate disclosure. High risk of discovery and use, even for modest potential harm, will accelerate disclosure.

Detect exploitation by others

How likely is it that we would know if someone else was exploiting it?

The government apparently has some capability: ‘After the discovery, the NSA tuned its sensors to detect use of any of the tools by other parties, especially foreign adversaries with strong cyber espionage operations, such as China and Russia.’

Effect. High confidence in ability to detect exploitation by others will delay disclosure; lower confidence moves disclosure forward.

Is the vulnerability needed?

How badly do we need the intelligence we think we can get from exploiting the vulnerability? Are there other ways we can get it?

This factor is essentially the government's own estimation of the value of access to a device and the information it contains. If there are other vulnerabilities than can be exploited—or other methods entirely—with less cost or risk, then those other methods might be preferable.

Effect. The existence of other methods of obtaining the desired information/access will reduce the value of retaining this vulnerability, and accelerate disclosure.

Discovery by others

How likely is it that someone else will discover the vulnerability?

Government concept of NOBUS: “that’s a vulnerability we are not ethically or legally compelled to try to patch — it’s one that ethically and legally we could try to exploit in order to keep Americans safe from others”

But recent studies have shown that simultaneous discovery of vulnerabilities happens.

Effect. If the vulnerability is likely to be discovered by others then it will accelerate disclosure. However, government confidence in a unique ability to discover or exploit some vulnerabilities will delay disclosure.

Can the vulnerability be used?

Could we utilize the vulnerability for a short period of time before we disclose it?

A few possible interpretations of this factor:

- it may simply not be possible to develop an exploit for a particular vulnerability
- Development could take too long
- Systems that could be accessed have no value

Effect. If the vulnerability cannot be utilized, then this will accelerate disclosure.

Can the vulnerability be patched?

Can the vulnerability be patched or otherwise mitigated?

Some systems such as SCADA, PLC, or embedded devices might not be able to be patched, although probably most can be mitigated.

Some systems are out of support, and patches will not be released for them. Disclosure of vulnerabilities might not be useful for users of these systems.

Effect. If patch creation and deployment can happen quickly, this will delay disclosure. If patching or mitigating is not possible, this will also delay disclosure.

Model

The model

$$V_{T*} = \max_{T} \left(B_0 + \sum_{t=1}^{T} d_b^t \, \mathbb{E}\left[B_t\right] - \left(C_0 + \sum_{t=1}^{T} d_c^t \, \mathbb{E}\left[C_t\right] \right) \right).$$

B_t and C_t are functions of the various factors.

d_b and d_c are discount factors. The process is supposed to be biased towards responsible disclosure — this could be reflected in the discount factors.

Effect on timing

Factor	F_{extent}	F_{harm}	F_{detect}	F_{value}	$F_{discovery}$	F_{use}	F_{patch}
Benefits	+			+		+	
Costs	+	+	-		+		-
Timing	-?	-	+	+	-	+	+

Influence of factors on the costs and benefits, compared to immediate disclosure, and how they affect the timing of disclosure. While F_{extent} influences both benefits and costs, it will likely have a greater influence on costs, moving disclosure forward.

EternalBlue and WannaCry

WannaCry Malware

Malware based on an NSA-developed exploit known as EternalBlue. Leaked to the public by the ShadowBrokers in **April** 2017.

Patch released by Microsoft in **March** — presumably disclosed by NSA.

Caused a lot of damage, affecting businesses and the UK NHS, and potentially could have caused a lot more. Only stopped because it contained a mechanism to stop when a particular domain name was registered.

How can we interpret this?

Interpretation 1

The decision was made using a correct model.

This implies that

- the vulnerability was disclosed at the appropriate time
- the benefits gained from the long-term retention of the vulnerability were valuable enough that they where not outweighed by the damages and costs that arose from the leaked vulnerability and resulting malware

Interpretation 2

The timing of the disclosure was wrong because the model was missing a factor: *the possibility of a vulnerability being leaked*.

From the Daniel blog post, we know that the risk of independent discovery is considered when making a decision, but it is unknown if this also includes the risk of leaks.

If not, then the time of disclosure would have been after the optimal point.

Interpretation 3

The timing of the disclosure was wrong because the model's parameters were incorrect.

Extent of use and patching factors:

- Patch was releaed by Microsoft before WannaCry, but many systems still vulnerable
- The rate at which patches can be developed and applied could be overestimated
- The number of out-of-support system could be underestimated
- Which lead to underestimation of cost, delaying disclosure

Incorrectly underestimating the probability of a leak (possibly included in the discovery factor) would also cause such a delay in disclosure.

Thanks!

Questions?

Tristan Caulfield

t.caulfield@ucl.ac.uk

@tristanc