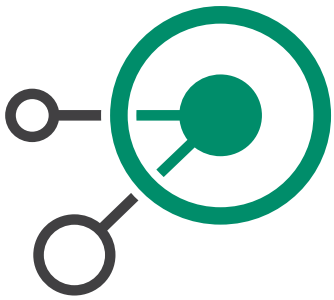

ERNW
providing security.

ERNW Newsletter

Reflections on Vulnerability Disclosure and a Case Study

ERNW
providing security.

Table of Content

List of Figures

List of Tables

1 Introduction & Background

Vulnerability disclosure has been a topic of fierce debates in the recent years. That's not in the least, usually a number of ethical questions are involved and proponents of different perspectives assign different weights and priorities to the values touched. In this paper we will discuss some of the questions involved, how they can be tackled and how we handle some of them in the past (and which developments make us consider it necessary to re-think our way of handling). The piece is organized as follows: first we provide a short overview of approaches to vulnerability disclosure and why we followed a specific one ("responsible disclosure"). We will then discuss potential problems with *responsible disclosure* which have arisen in the interim. To illustrate these (types of) issues we will discuss a specific case study we've been involved with. Furthermore, we will formulate a set of questions to stimulate further discussion of the topic. It should be noted that this paper is written from a highly personal perspective and it's not meant to provide definitive answers, but to raise awareness of the inherent challenges of the process.

1.1 The Role of Security Research at ERNW

Since our first days as a company (back in 2001) we have always been performing substantial security research at ERNW, for several reasons:

- to develop our capabilities/skills and our methodology when facing certain tasks.
- to contribute to public security knowledge & discussion (or, as we put it at Troopers: "to make the world a safer place"). The people running the organization, including myself, consider it an important duty of any company to contribute to society in general (by paying taxes, providing employment, developing talent etc.) and particularly to commit to the field of expertise one is working in.
- it helps to increase the visibility of our expertise which in turn helps to achieve "economic sustainability of the company as a whole" (which, let me be clear here, is probably amongst the objectives of any reasonable company).
- simply because security research is fun ;-).

1.2 Types of (Research) Projects Leading to the Discovery of Vulnerabilities

From a "project sponsor perspective" the research activities we undertake can be grouped into three main categories:

- research we "just do on our own" (read: without a specific customer context), because we think it's important to look at the security properties of a class of devices or because we're curious as for the real-life implementation of protocols etc. Actually every ERNW member (except for the back office staff) is expected to participate in at least one ongoing activity of this type.
- research that is somewhat related to/sponsored by a customer security assessment project. Here we often come up with an agreement along the lines of "while you [customer] pay n man-days for the assessment, we're willing to spend much more effort for a certain component if you're ok that we share the results with the public thereafter (of course, without any reference to the specific environment)"[1].
- research projects we're engaged for in a dedicated manner. The main property here (at least in the context of the present essay) being that the engaging party fully owns the intellectual property from the project so we're not necessarily involved in the way the results are disseminated in the end of the day.

In this paper we cover only first two categories (for the simple reason we have some decision power there) and the inherent question: what to do with vulnerabilities found in the course of those?

[1] *Examples include* https://www.ernw.de/download/BlackHat-EU-2010-Mende-Rey-Cisco-Wlansec-slides.pdf *or https://www.insinuator.net/2014/11/github-enterprise-2-0-0-fixes-multiple-vulnerabilities/.*

1.3 Vulnerability Disclosure – Main Approaches & Some Additional Sources

From a simplistic perspective considering the "audience" in the early stages of the procedure the following variants of vulnerability disclosure can be distinguished:

- tell everybody – this is usually called "full disclosure".
- initially only tell the vendor (of a product) and release information to the public at a later point – this is usually called "responsible disclosure"[2]. An important variable here is the time period imposed by the researcher.
- do something else, which includes telling a broker, telling some 3rd party (usually in exchange for money), telling nobody or any other possible course of action.

For those interested in specific flavors, their real-life shapes or their advantages/disadvantages the following sources can be useful[3]:

- Details on Google's Project Zero policy: *http://googleprojectzero.blogspot.de/2015/02/feedback-and-data-driven-updates-to.html*.
- Approach of the Carnegie Mellon CERT: *https://www.cert.org/vulnerability-analysis/vul-disclosure.cfm*.
- Paper "Emerging Economic Models for Vulnerability Research" from the 2006 *Workshop on the Economics of Information Security*: *http://weis2006.econinfosec.org/docs/17.pdf*.
- There's an ISO standard: ISO/IEC 29147:2014 "Information technology -- Security techniques -- Vulnerability disclosure" (focusing on the vendor perspective).
- In 2002 there was even an IETF draft: https://tools.ietf.org/id/draft-christey-wysopal-vuln-disclosure-00.txt.
- The *US Department of Commerce – National Telecommunications & Information Administration* just recently initiated a "multistakeholder process" to discuss the topic: *http://www.ntia.doc.gov/other-publication/2015/multistakeholder-process-cybersecurity-vulnerabilities*.
- Three Italian researchers provided their positions in response to a blog post I wrote on the topic: *http://blog.nibblesec.org/2015/08/vulnerability-disclosure-what-could_17.html*.

Our own approach has been significantly shaped by the so-called RFPolicy[4] which stated, amongst others:

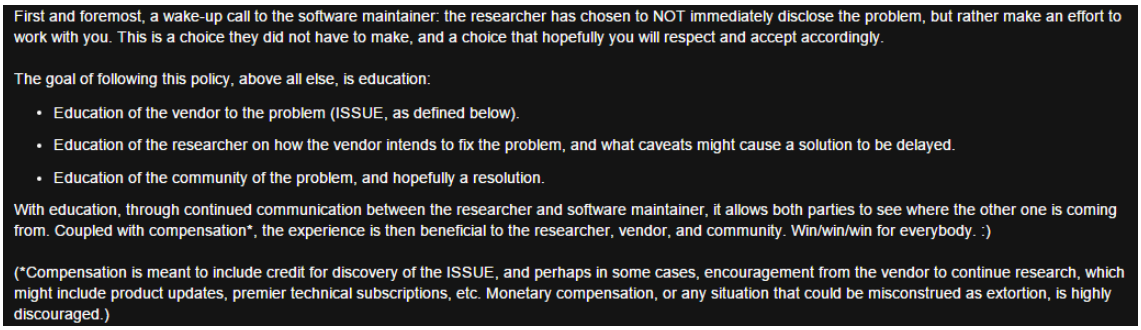
First and foremost, a wake-up call to the software maintainer: the researcher has chosen to NOT immediately disclose the problem, but rather make an effort to work with you. This is a choice they did not have to make, and a choice that hopefully you will respect and accept accordingly.

The goal of following this policy, above all else, is education:

- Education of the vendor to the problem (ISSUE, as defined below).
- Education of the researcher on how the vendor intends to fix the problem, and what caveats might cause a solution to be delayed.
- Education of the community of the problem, and hopefully a resolution.

With education, through continued communication between the researcher and software maintainer, it allows both parties to see where the other one is coming from. Coupled with compensation*, the experience is then beneficial to the researcher, vendor, and community. Win/win/win for everybody. :)

(*Compensation is meant to include credit for discovery of the ISSUE, and perhaps in some cases, encouragement from the vendor to continue research, which might include product updates, premier technical subscriptions, etc. Monetary compensation, or any situation that could be misconstrued as extortion, is highly discouraged.)

Figure 1: Excerpt from RFPolicy

[2] Microsoft has meanwhile coined the term "Coordinated Disclosure" for this approach, see *https://technet.microsoft.com/en-us/security/dn467923.aspx*.

[3] *Probably the most comprehensive overview of the history of the process is Haroon Meer's compilation:* https://www.duosecurity.com/labs/timelines/history-of-the-disclosure, *although it has been pointed out by Ivan Arce that some important elements are missing there.*

[4] *See* http://www.wiretrip.net/p/rfpolicy.html.

2 A MORE DETAILED DISCUSSION OF *RESPONSIBLE DISCLOSURE*

In the simplest scenario (and using ISO 29147 terminology) there's two actors involved:

- the *finder*[5] who has discovered a vulnerability which they now report
- the *vendor*[6] who receives the information
- and subsequently provides *remediation*[7].

Referring to the *RFPolicy* the overall objective of this approach can be summarized as follows: "contribute to the education of the parties involved/affected and thereby help to achieve an overall higher state of security for everybody". We designate this objective as [OBJ_L_PUBLIC_CULTURE], where "L" stands for "long-term".

2.1 Inherent Assumptions

Usually the following inherent assumptions apply:

- there's no patch available at the time of reporting.
- the vendor actually takes care of remediation.
- once this remediation is public (e.g. a security advisory and/or a patch is released), it can be deployed everywhere where needed, without too much delay.
- The people involved/users affected (let's call them the "stakeholders") are well-informed, willing to deploy the remediation and enabled (with regard to their technical skills and the environmental conditions) to do so.

You might already note that this is a quite greenfield scenario and accompanying set of assumptions. Still, going with a fairly standard responsible disclosure approach (depending on the specific conditions and overall picture with a 30 days or a 90 days period) has worked for us in the vast majority of cases.
For us taking the path of a *RFPolicy*-oriented approach also implied that in the course of each responsible disclosure process we went through:

- we always got into direct contact with the vendors. So far we've never went through brokering organizations like HP/TippingPoint's Zero Day Initiative[8].
- we never asked for or received any financial compensation, neither in a direct way ("here's some money") nor in an indirect way ("what about bringing you guys in as a new pentesting partner? [of course we need to establish a good trust-based relationship first and you, ERNW, should start building this by refraining from any publication]...").
- we have never sold any vulnerability information to a 3rd party.

The above is not about being moralistic and we're not judging other ways of handling. It's just, as I said, that this scheme of things worked for us and we considered it a good policy to act in a way consistent with our values and objectives.

[5] ISO 29147, sect. 3.3 "finder: individual or organization that identifies a potential vulnerability in a product or online service".
[6] ISO 29147, sect. 3.8 „vendor: individual or organization that developed the product or service or is responsible for maintaining it".
[7] ISO 29147, sect. 3.6 "remediation: patch, fix, upgrade, configuration, or documentation change to either remove or mitigate a vulnerability".
[8] http://www.zerodayinitiative.com/about/

2.2 Potential Problems

Now let's have a look at some developments we observe that make us reconsider if this is still the right approach. Here, for a moment, let's keep in mind that my duties as managing director of ERNW include:

- taking care that our actions are conformable to the law.
- taking care that our actions are consistent with our values and our derived overall objectives.
- taking care that our resources, incl. my own, are spent in a way contributing to our objectives.

In the context of (responsible) vulnerability disclosure the above becomes increasingly difficult, mainly for two reasons outlined in the following sections.

2.2.1 Legal Blur

I have the impression that there is – compared with, say: 3-5 years ago – a growing number of vendors out there which operate with outspoken or elusive legal threats in the course of the procedure (which is a bit surprising as for some time it seemed that responsible disclosure had become a well established practice, not in the least due to the efforts people like Katie Moussouris spent on developing ISO 29147, ISO 30111 and the like). Personally I have to say I'm increasingly tired of this. We have responsibly reported numerous bugs which due to their nature would have raised significant sums of money if reported through a broker or sold "to interested 3rd parties". I hence feel less and less inclined to go through series of conference calls to hear cryptic mentions of the "our legal department is working on this, to protect our customers" type, and to constantly remind myself: "it's the right thing to do, for the long-term sake of the community of stakeholders".

Furthermore, there is the *Wassenaar Arrangement* (WA). I won't go into a detailed discussion here as many people smarter than me have expressed their views (e.g. Sergey Bratus[9]) and there's extensive and enlightening discussions on the "Regs — Discussions on Wassenaar" mailing list[10] established by Arrigo Triulzi. Still suffice to say here that it is my understanding that the Wassenaar Arrangement has severe implications with regard to the way vulnerability disclosure, once performed "across borders", takes place. For example, we're right now involved in a disclosure procedure (on a nasty chain of bugs identified by a group of researchers) with a vendor of security appliances. Let's assume that the vendor is located in another country than Germany. How are we supposed to provide Proof-of-Concept code to the vendor in a timely manner once such code could be considered to be covered by the second, controlled class of software as of the WA? We don't think it is (covered), but we're not keen on going into lengthy legal battles over this estimation either. Those interested in the practical implications might have a look at the case of Grant Willcox, a British researcher who initially removed some parts of the public version of his final year's dissertation ("An Evaluation of the Effectiveness of EMET 5.1") due to the WA[11].

Next to this type of concerns there's another main aspect which keeps us reflecting on the suitability of responsible disclosure. From a high-level perspective as for the overall objective of responsible disclosure this one could be even more important.

[9] http://www.cs.dartmouth.edu/~sergey/drafts/wassenaar-public-comment.pdf.
[10] https://lists.alchemistowl.org/mailman/listinfo/regs/.
[11] *http://tekwizz123.blogspot.de/2015/07/final-year-dissertation-paper-release.html.*

2.2.2 New Stakeholders & Objectives

As outlined above, the RFPolicy inspired approach included two main classes of actors (the finder and the vendor), with an additional vague mention of "the community", and it worked on the basis of some (inherent) assumptions. However nowadays a number of cases we're involved with are quite different in one way or another. The main differences that we can identify are the following:

- increasingly there's another group of stakeholders involved which are not part of the above, "traditional" picture, but who are heavily affected (when driving their car, when being treated by means of network-connected medical devices, when using some piece of technology in their household or even using pieces of technology to protect this very sphere etc.).
- the vulnerabilities in question might have a direct impact on their health or on their personal property (as opposed to the somewhat anonymous assets of enterprise organizations or vendors depicted in the classic RFPolicy).
- at the same time the affected users might be completely unaware of the vulnerabilities.
- even if they knew, due to the specific nature of certain components/devices it might just not be technically possible or feasible to apply the remediation.

These aspects induce another main objective (of vulnerability handling), to be designated as follows:

[OBJ_S_PUBLIC_PREV_HARM]: "protect public from harm against their lives, health or economic situation" (where the "S" marks that this usually is a – somewhat – short-term goal).

Identifying this objective evidently brings up an interesting question: how to proceed once the now two main objectives (of vulnerability [non-] disclosure), that are [OBJ_L_PUBLIC_CULTURE] and [OBJ_S_PUBLIC_PREV_HARM], clash? Or, to put it less abstract: what if pursuing the long-term goal of (vendor/community) education conflicts with the short-term goal of not contributing to people getting harmed?

I'll provide a simple example: are we supposed (or even obliged from an ethical point of view) to disclose vulnerabilities in a medical device (maybe, after having tried to get in contact with the vendor several times and on several channels, without luck)? Public disclosure might put patients in danger (and the devices possibly can't be patched anyway, for regulatory reasons). On the other hand: whom does it help if we just sit on the information? Should we try to go through other channels? If so, which ones?

As I stated in the introduction we don't have an easy answer for this type of situations with conflicting objectives. Right now we handle those on a case-by-case basis (see the case study below or our research results with regard to certain alarm systems[12]). We have even established an ethics committee at ERNW about two years ago, not least in order to resolve this type of dilemmas. It can be consulted by every member and then the committee is entitled to provide a recommendation considered binding for everybody, including management.

Still we keep thinking there might be better/more suitable ways of vulnerability handling (and there's probably several other researchers facing the same type of questions).

[12] https://www.insinuator.net/2015/05/analysis-of-an-alarm-system-part-23/. *In this case we refrained from publishing the third part of that series so far.*

3 Conclusions & Outlook

Given we've not yet come up with a satisfying let's have a quick look at possible alternatives. These include:

- don't do anything with vulnerabilities we discover and "just sit on them", maybe for a certain period of time imposed by some governing rules we have to come up with, maybe "indefinitely".
- go full disclosure.
- go through a broker (which saves energy & time, too. urthermore, this could bring in money to be used for additional Troopers student invitations, the Troopers charity fund or just some more nice equipment for the lab. I'm sure the guys would come up with plenty of ideas...).
- only report to vendor once there's a bug bounty program (alternatively "drop 0day" as our old buddy Michael Ossmann suggested[13]).
- perform full disclosure and combine it with going through media/the press (again this could save energy & time and it might even increase the reach, hence subsequently contribute to the objective of "public education").
- hand over everything to something like a "national clearing house".
- something else...

For the moment we don't find any of those particularly consistent with the overall objectives. Still we sense we have to develop an adapted approach to vulnerability disclosure, for the reasons outlined above. It's just: what could that new approach look like? We leave that as an exercise to the reader ;-) and we're happy to receive any type of feedback.

[13] *https://twitter.com/michaelossmann/status/595417174958841856.*

4 Case Study

In this chapter we describe a case study which illustrates some of the issues discussed above. We will describe some events and the associated (ethical) questions they raised for us. Again, this is mostly meant to stimulate thinking through this topic, as we don't feel qualified to provide definite answers or guidelines, and – in hindsight – we're somewhat unhappy with certain decisions we took. In any case we hope that reading this might help others who face the same issues to come up with decisions more easily.

4.1 Overview

Mainstream media reported (initially on Jan 28 2014) that the most popular residential gateway product used in Germany, the so-called "Fritz!Box"[14] was abused in the wild, in a way that generated significant telephone bills for the users affected. Some days later (on Feb 3 2014) the prominent German IT news outlet "heise online" (in the following just "heise") reported on the issue and laid out that probably the remote access ("Fernzugriff") capability was involved. In the same week (on Feb 6 2014) AVM released a security advisory explaining that the remote access capability was actually involved but that this function was disabled by default (which means users would be unaffected as long as they had not explicitly enabled it). Further they recommended disabling the capability where enabled, and this advice was echoed by *heise*.

The vendor of the devices ("AVM Computersysteme Vertriebs GmbH", in the following just "AVM") then investigated these cases and very quickly (that is starting from Feb 8 2014) provided firmware updates (hence "patches") for some popular models[15].

Shortly thereafter, an ERNW member took a closer look at a patch[16] and determined that actually the remote access capability (deemed disabled by default) was *not* responsible for the vulnerability, but another function was involved and that the bug could be exploited just by a user sitting behind a vulnerable device visiting a malicious website. This meant that a much larger user base was potentially affected and that the media (and their audience's) perception of not being vulnerable as long as the remote access capability was not willingly enabled was wrong. This created a situation which heavily differed from the vast majority of vulnerability disclosure scenarios we had handled in the past and which did not fulfill many of the inherent assumptions laid out above. It was further complicated by the following aspects:

- the vulnerability was apparently exploited in the wild.
- a patch was actually available, but couldn't be rolled out fast enough.
- it had to be assumed that many affected users unaware of problem and could not be easily reached (not least because many devices were deployed as CPEs via ISPs).
- mainstream media was (or quickly became) aware of the issue and had their own agenda, potentially differing from the agendas of the other players involved.
- it was not about a "0-day" vulnerability but about a widely discussed one which could be determined by simply looking at the patch with "reasonable effort".
- some media (continued to) spread the false information that "remote management would have to be enabled" to exploit the vulnerability.

A timeline containing a number of things that happened can be found on the next page.

---

[14] *See also* http://en.wikipedia.org/wiki/Fritz!Box.

[15] *It should be noted that the rapid provision/making available of patches was – rightfully, as of the authors of the present paper – considered as a huge and very laudable undertaking from the vendor.*

[16] *Anecdote: he did so after he was approached by a family member: "I heard there's an issue with the Fritzboxes. Can you, as a hacker, tell me what's going on?"*

Number	Date	Event
1	Jan 28 2014	Initial news[17] on abuse/toll fraud taking place and potentially related to devices.
2	Feb 3 2014	First report on heise.de[18], speculating that remote management feature involved.
3	Feb 6 2014	AVM releases bulletin laying out remote management capability involved.[19]
4	Feb 7 2014	Heise writes that remote mgmt. involved and recommends to disable it[20].
5	Feb 8 2014	AVM (impressingly quickly!) starts releasing firmware updates/patches.
6	Feb 8 2014	Analysis of patch shows that vulnerability much more severe than expected and that it can be exploited *without* remote mgmt feature enabled.
7	Feb 9 2014	Heise announces that AVM has released patches[21].
8	Feb 10 2014	Mainstream media[22] pick up the story but with misleading description.
9	Feb 10 2014	We establish contact to BSI to discuss details and we try to reach out to vendor.
10	Feb 11 2014	Vendors call us and asks us to not disclose details to prevent further/copycat attacks. They told us there were about eight million devices in the field and they could patch 500.000 per day (=> in best case 16 days needed if all could be reached which, for several reasons, was unrealistic/wouldn't happen anyway).
11	Feb 11 2014	Some law enforcement agency calls us and asks us not to disclose details as copycat attacks would impact ongoing investigation. They did not issue any threats; actually this was a very friendly and cooperative call.
12	Feb 11 2014	BSI issues press release[23]. This did not mention different attack path but stresses importance of deploying the patch.
13	Feb 17 2014	"heise Security" publishes report that remote management does not have to be enabled to exploit the vulnerability[24]. This included a video demonstrating the attack.
14	Feb 18 2014	Attack & impact discussed in Germany's most prominent news broadcast "Tagesschau".
15	Mar 7 2014	Heise reports that details of the attack are publicly available[25].
16	Mar 11 2014	We publish our post[26] (which was meant to be a mostly educational piece anyway).

Table 1: Timeline of Events in AVM Case

[17] http://www.derwesten.de/staedte/nachrichten-aus-moers-kamp-lintfort-neukirchen-vluyn-rheurdt-und-issum/hacker-greifen-telefonkunden-an-id8922516.html.
[18] http://www.heise.de/security/meldung/Telefonie-Missbrauch-anscheinend-kein-Massenhack-von-AVMs-Fritzboxen-2104609.html.
[19] http://avm.de/aktuelles/kurz-notiert/2014/wichtiger-sicherheitshinweis-fuer-fritzbox-nutzer-mit-aktiviertem-fernzugriff/.
[20] http://www.heise.de/security/meldung/Fritzbox-Angriff-analysiert-AVM-bietet-erste-Firmware-Updates-an-2108862.html
[21] http://www.heise.de/security/meldung/Fritzbox-Hack-AVM-bietet-Updates-fuer-ueber-30-Modelle-2109372.html
[22] http://www.spiegel.de/netzwelt/web/fritzbox-router-hersteller-avm-veroeffentlicht-sicherheits-updates-a-952468.html
[23] https://www.bsi.bund.de/DE/Presse/Pressemitteilungen/Presse2014/Fritz-Box-Update_11022014.html
[24] http://www.heise.de/security/meldung/Jetzt-Fritzbox-aktualisieren-Hack-gegen-AVM-Router-auch-ohne-Fernzugang-2115745.html.
[25] *http://www.heise.de/security/meldung/Hack-gegen-AVM-Router-Fritzbox-Luecke-offengelegt-Millionen-Router-in-Gefahr-2136784.html.*
[26] http://www.insinuator.net/2014/03/how-to-own-a-router-fritzbox-avm-vulnerability-analysis/.

4.2 Objectives & Questions in the Individual Phases

In this section I'd like to present some of the questions we were facing at the time.

4.2.1 Objectives in Initial Phase

From today's perspective, with a structured approach, the following objectives influencing how to handle the situation initially can be identified:

- [OBJ_S_PUBLIC_PREV_HARM]: protect public (users/households running devices without evening knowing) from financial harm by high telephone bills (in the range of several thousand EUR).
- [OBJ_S_PUBLIC_AWARENESS]: create awareness that problem exists and that misconceptions might be involved ("devices only vulnerable with remote mgmt. enabled").
- [OBJ_L_PUBLIC_CULTURE]: the "traditional" long-term objective to contribute to greater good by applying a certain disclosure approach, as laid out in the above standard responsible disclosure setting.
- [FINDER_RECOGNITION]: evidently in the overall picture the *finder* is an actor as well and as such has an agenda. In this case we would have been happy to be visible both as "guys with technical expertise" and "guys acting in a certain [maybe responsible, maybe not] way".

Main Questions We Asked Ourselves

- Is "responsible disclosure" the right approach to handle the situation? [given vendor knows already about the problem, patches are available, and exploitation happens in the wild]

What We Actually Did

Traditionally, in responsible disclosure procedures, [OBJ_L_PUBLIC_CULTURE] has been somewhat prioritized over [OBJ_S_PUBLIC_PREV_HARM] – not least because vendors practically always request not to disclose citing the exact latter one… – but in this case we decided to follow a different path and contacted the German Federal Office for Information Security ("Bundesamt für Sicherheit in der Informationstechnik", BSI[27]) which has a department responsible for protection and awareness of citizens/the public. In parallel (on Feb 10 2014) the vendor was contacted (but they did not react until contacted from the BSI).

In Hindsight

We have not come up with better ideas.

4.2.2 After Being Contacted from Law Enforcement Agency ("Event 11")

Main Questions We Asked Ourselves

- Does this change the situation/priorities?

What We Actually Did

We confirmed that we wouldn't publish without advance notification and we added another party to the picture whose (very legitimate) concerns would have to be taken into account when it came to taking decisions.

In Hindsight

We have not come up with another way of handling.

[27] https://www.bsi.bund.de/EN/TheBSI/thebsi_node.html.

4.2.3 After Heise Published the Video and Made it Clear that the Attack Path Was Different ("Event 13")

Main Questions We Asked Ourselves

- Does this change the situation/priorities?
- Worse or better (for the stakeholders)?
- Does this help reaching (anybody's) objectives?

What We Actually Did

We waited/did not publish as we knew the patches technically could not yet be deployed "in large numbers". Still, from a [FINDER_RECOGNITION] perspective, this was somewhat unfortunate and, more importantly, we did not think that the Heise article/video helped anybody. But, sure, different parties & different agendas…
Furthermore, we had given our word to both the CTO of the vendor and to our contact person at the law enforcement agency that we would not publish without "reasonable advance notification" so we couldn't have reacted quickly anyway.

In Hindsight

That's an interesting question. Actually I don't think it would have had an impact on [OBJ_S_PUBLIC_PREV_HARM] if we had published ("bad actors knew the stuff anyway, latest after that video, 'even if [that one was] obfuscated'").
Probably it wouldn't have changed [OBJ_S_PUBLIC_AWARENESS] either (realistically our blog doesn't have much reach outside the technical community). So, (solely) looking at the above objectives, there was no real benefit from our (non-) action.

4.2.4 After Heise Announced that Details and Exploit Code Are Easily Available ("Event 15")

Main Questions We Asked Ourselves

- Does this change the situation/priorities?
- Worse or better (for the stakeholders)?
- Does this help reaching (anybody's) objectives?

What We Actually Did

We informed our points of contacts at different parties that we were going to publish then ("the cat is out of the bag now") and we waited some additional days.

4.2.5 Today's Feelings

Overall I doubt that withholding the post for four weeks "contributed to the greater good" or really helped to achieve [OBJ_S_PUBLIC_PREV_HARM] or [OBJ_S_PUBLIC_AWARENESS]. In addition, (not publishing) probably it did not add much value to [OBJ_L_PUBLIC_CULTURE].

On the other hand we felt some sympathy for the vendor who did really good efforts, and at least this course of action gave us some warm, cozy feeling of "not being directly responsible that households were getting ripped".

4.3 Lessons Learned & Recommendations

Which conclusions can be derived from the above case that could be worthwhile for other researchers facing similar dilemmas? Here's some advice we'd like to provide:

- Try to clearly identify the stakeholders involved (incl. "secondary actors" like the press or law enforcement entities).
- Make up your mind with regard to the objectives of your way of handling. Which of those mentioned in section 5.2.1 or which different ones do you want to support by your approach? Who are the beneficiaries of different strategies?
- Prioritize objectives accordingly. Usually you can't fully reach them all; there will be conflicts.
- Try to distinguish between short/medium/long-term goals. Are regulatory requirements kicking in?
- Constantly re-evaluate the above, in particular in case new developments happen.