Google Security Blog

The latest news and insights from Google on security and safety on the Internet

Rebooting Responsible Disclosure: a focus on protecting end users

July 20, 2010

Posted by Chris Evans, Eric Grosse, Neel Mehta, Matt Moore, Tavis Ormandy, Julien Tinnes, Michal Zalewski; Google Security Team

Vulnerability disclosure policies have become a hot topic in recent years. Security researchers generally practice “responsible disclosure”, which involves privately notifying affected software vendors of vulnerabilities. The vendors then typically address the vulnerability at some later date, and the researcher reveals full details publicly at or after this time. A competing philosophy, "full disclosure", involves the researcher making full details of a vulnerability available to everybody simultaneously, giving no preferential treatment to any single party. The argument for responsible disclosure goes briefly thus: by giving the vendor the chance to patch the vulnerability before details are public, end users of the affected software are not put at undue risk, and are safer. Conversely, the argument for full disclosure proceeds: because a given bug may be under active exploitation, full disclosure enables immediate preventative action, and pressures vendors for fast fixes. Speedy fixes, in turn, make users safer by reducing the number of vulnerabilities available to attackers at any given time. Note that there's no particular consensus on which disclosure policy is safer for users. Although responsible disclosure is more common, we recommend this 2001 post by Bruce Schneier as background reading

on some of the advantages and disadvantages of both approaches. So, is the current take on responsible disclosure working to best protect end users in 2010? Not in all cases, no. The emotionally loaded name suggests that it is the most responsible way to conduct vulnerability research - but if we define being responsible as doing whatever it best takes to make end users safer, we will find a disconnect. We've seen an increase in vendors invoking the principles of "responsible" disclosure to delay fixing vulnerabilities indefinitely, sometimes for years; in that timeframe, these flaws are often rediscovered and used by rogue parties using the same tools and methodologies used by ethical researchers. The important implication of referring to this process as "responsible" is that researchers who do not comply are seen as behaving improperly. However, the inverse situation is often true: it can be irresponsible to permit a flaw to remain live for such an extended period of time. Skilled attackers are using 0-day vulnerabilities in the wild, and there are increasing instances of:

- 0-day attacks that rely on vulnerabilities known to the vendor for a long while.
- Situations where it became clear that a vulnerability was being actively exploited in the wild, subsequent to the bug being fixed or disclosed.

Accordingly, we believe that responsible disclosure is a two-way street. Vendors, as well as researchers, must act responsibly. Serious bugs should be fixed within a reasonable timescale. Whilst every bug is unique, we would suggest that 60 days is a reasonable upper bound for a genuinely critical issue in widely deployed software. This time scale is only meant to apply to critical issues. Some bugs are mischaracterized as "critical", but we look to established guidelines to help make these important distinctions — e.g. Chromium severity guidelines and Mozilla severity ratings. As software engineers, we understand the pain of trying to fix, test and release a product rapidly; this especially applies to widely-deployed and complicated client software. Recognizing this, we put a lot of effort into keeping our release processes agile so that security fixes can be pushed out to users as quickly as possible. A lot of talented security researchers work at Google. These researchers discover many vulnerabilities in products from vendors across the

board, and they share a detailed analysis of their findings with vendors to help them get started on patch development. We will be supportive of the following practices by our researchers:

- Placing a disclosure deadline on any serious vulnerability they report, consistent with complexity of the fix. (For example, a design error needs more time to address than a simple memory corruption bug).
- Responding to a missed disclosure deadline or refusal to address the problem by publishing an analysis of the vulnerability, along with any suggested workarounds.
- Setting an aggressive disclosure deadline where there exists evidence that blackhats already have knowledge of a given bug.

We of course expect to be held to the same standards ourselves. We recognize that we've handled bug reports in the past where we've been unable to meet reasonable publication deadlines -- due to unexpected dependencies, code complexity, or even simple mix-ups. In other instances, we've simply disagreed with a researcher on the scope or severity of a bug. In all these above cases, we've been happy for publication to proceed, and grateful for the heads-up. We would invite other researchers to join us in using the proposed disclosure deadlines to drive faster security response efforts. Creating pressure towards more reasonably-timed fixes will result in smaller windows of opportunity for blackhats to abuse vulnerabilities. In our opinion, this small tweak to the rules of engagement will result in greater overall safety for users of the Internet.

Update September 10, 2010: We'd like to clarify a few of the points above about how we approach the issue of vulnerability disclosure. While we believe vendors have an obligation to be responsive, the 60 day period before public notification about critical bugs is not intended to be a punishment for unresponsive vendors. We understand that not all bugs can be fixed in 60 days, although many can and should be. Rather, we thought of 60 days when considering how large the window of exposure for a critical vulnerability should be permitted to grow before users are best served by hearing enough details to make a decision about implementing

possible mitigations, such as disabling a service, restricting access, setting a killbit, or contacting the vendor for more information. In most cases, we don't feel it's in people's best interest to be kept in the dark about critical vulnerabilities affecting their software for any longer period.

Update May 12, 2021: February 2015, Google's responsible discloure policy was adjusted to 90 days. View our current policy here.


29 comments :

Harry Johnston said...

It should be noted that the argument that full disclosure "enables immediate preventative action" is only true for a small subset of computers - the administrator has to be both knowledgeable enough to be aware of the problem and able to make configuration changes straight away, i.e., without testing them first.

Most home users don't follow full-disclosure mailing lists, and in many corporate scenarios extended testing is necessary to check for any possible impact to line-of-business applications before a change can be made.

Also, I'm doubtful that 60 days is long enough to properly test an update to ensure it won't break anything for end users. In the absence of evidence that a specific vulnerability is already being abused, is it really necessary for vendors to be so hasty in releasing an update?

July 20, 2010 at 8:24 PM

Justin said...

60 days is far too generous for waiting for full disclosure. Unless it takes a complete rewrite of the program, what takes that long to fix and test?

July 20, 2010 at 9:30 PM

Harry Johnston said...

To pick a specific example not entirely at random, let's talk about Microsoft Windows. :-)

Windows (like all other modern general-purpose operating systems) is a very, very complicated piece of software. It takes time to properly test any proposed change, along with the associated update mechanism and so on. Don't forget that changes have to be tested for each supported variant of the OS, as well as on a variety of hardware platforms.

Despite the extensive testing Microsoft already do, I've seen more than one person complain - with good justification - that Windows updates aren't reliable enough. If you look after hundreds or thousands of computers, the last thing you want is to apply an update that was rushed out the door, or, for that matter, a quick-fix workaround.

If the black hats discovered a vulnerability themselves, well, that's just the way things go. But to be forced to take risks because a researcher didn't want to wait until a reliable fix was ready is frustrating, to the say the least.

July 20, 2010 at 10:30 PM

Harry Johnston said...

Oh, and I forgot to mention: don't forget that the vendor may be working on more than one vulnerability at a time.

July 20, 2010 at 11:45 PM

Eliot Lear said...

Google is well known for being a data analysis company that happens to offer a search engine, and so I wouldn't dispute the claim that 0-day attacks are on the rise. There is also other research to support such a claim (certainly SANS said so in November of last year).

However, this article shows a narrow view of where vulnerabilities can occur. If a vulnerability is in a browser on a general purpose computer, one would think that 60 days is plenty time to correct a problem. On the other hand, if a vulnerability is in a core component of critical infrastructure that requires re-certification after any changes, that process can take quite a while. Some governments are in fact calling for more of such certifications in an attempt to improve interoperability. Interesting how such certifications can actually work against security.

Even in the case of a general purpose computer, the ability to patch that computer will be limited by its purpose. If it controls a power plant or an MRI system, perhaps more care will be given than if it sits on my dad's desk.

A substantial amount of research has occurred in this area. Readers and Google might wish to review past research that has been presented at the Workshop on the Economics of Information Security (WEIS) on optimal patch strategies. One such paper (found by Google!) is **Information Security Trade-offs and Optimal Patching Policies** by Ioannidis, Pym, and Willaims.

July 21, 2010 at 1:51 AM

Anonymous said...

What Tavis Ormandy did was cyber terrorsim, stop trying to defend him via this blog.

July 21, 2010 at 5:16 AM

B said...

Quoting:
"We of course expect to be held to the same standards ourselves. We recognize that we've handled bug reports in the past where we've been unable to meet reasonable publication deadlines -- due to unexpected dependencies, code complexity, or even simple mix-ups. In other instances, we've simply disagreed with a researcher on the scope or severity of a bug. In all these above cases, we've been happy for publication to proceed, and grateful for the heads-up."

--

I've experienced this "special treatment" by Google employees as well, let me tell you my story.

About three years ago I have found, reported and publicized srveral security flaws in Google Mail & Google Groups and later in Google's global gaia module used for authentication.

If I were a bad guy, these flaws would have enabled me to obtain access to all Google Services a specific user is using, including Adsense, Analytics and GMail.

I reported these flaws and what did I get in return?

- My GMail account was purged without warning.
- My Adsense / Analytics accounts were banned, and still my name seems to be on the registration prohibited blacklist
- basically all of my Accounts were deleted.

After years of sending emails to various support addresses I haven't received a single answer.

So long and thanks for nothing,

July 21, 2010 at 8:50 AM

Unknown said...

Eliot- Good luck trying to find data on 0day attacks, they're called "0day" attacks for a reason :P

Also what companies(besides google) are going to publish the fact that 0day attacks were used on them?

Sorry if it offends your analytical nature, but there is not going to be data on this...you wont be able to find ROI on this...this is hacking, get used to it.

60 days is plenty of time for a vendor that has any care in the world for security.

July 21, 2010 at 10:17 AM

Anonymous said...

Since your concept involves full disclosure after what you state (and many of us agree) is a reasonable grace period for vendor resolution, and since you take note of the power of labels like "responsible disclosure", I propose that you put a name to this philosophy of bug disclosure. My suggestion is simply "reasonable disclosure." At the very least, putting a label on it facilitates discussion of its merits and the word reasonable helps combat the emotional power you point out exists in "responsible."

July 21, 2010 at 10:37 AM

Harry Johnston said...

Quoting:
"We of course expect to be held to the same standards ourselves. [...] In all these above cases, we've been happy for publication to proceed, and grateful for the heads-up."

There's another issue hidden away here. It isn't really the vendor who is worst affected by the premature disclosure of a vulnerability. You may have been happy for publication to proceed, but what about your customers?

I foresee a lawsuit, one day, and the vendor won't be the plaintiff: that will be one of the vendor's customers, whose essential systems were disrupted by an attack based on a prematurely disclosed vulnerability. Ironically enough, it wouldn't be surprising if the vendor

and the security researcher(s) were listed as co-defendants.

(I'm guessing that such a lawsuit wouldn't fly in the US because it would be seen as a first amendment issue, but in other parts of the world it could be a real risk.)

July 21, 2010 at 2:34 PM

Harry Johnston said...

n3tv3d:

Aggravating as that incident was, nobody died. Politicians cheapen the word "terrorism" enough as it is, let's not make it any worse.

July 21, 2010 at 2:37 PM

Anonymous said...

Why do people keep associating cyber terrorism with people dying?

Cyber terrorism is nothing to do with people dying, cyber terrorism is a different type of terrorism than physical terrorism.

July 21, 2010 at 3:06 PM

Harry Johnston said...

n3td3v says: "Cyber terrorism is nothing to do with people dying, cyber terrorism is a different type of terrorism than physical terrorism."

You mean, the kind that isn't actually terrorism at all?

Misusing the word in this way is disrespectful to the victims of actual terrorist attacks.

July 21, 2010 at 5:15 PM

Anonymous said...

You just don't understand what *cyber* terrorism.

You're being clouded by visions of twin tower attacks by the sounds of it.

Cyber terrorism isn't about killing *people*.

It's about cyber attacks on against computers.

July 21, 2010 at 7:48 PM

newsoft said...

Disclosure policy has no effect whatsoever on software quality, therefore end-users protection.

Even if Adobe was able to fix Adobe Reader in 2 days, as long as product design and implementation were flawed in the first place, the product will remain a well of bugs that are waiting to be found (and exploited).

According to Secunia, this holds true for most consumer software products (e.g. IE6, Office 2003, Adobe Reader, Flash Player, Sun JVM - you know what I mean :), Apple QuickTime, RealPlayer, and so on).

July 22, 2010 at 2:51 AM

Unknown said...

as a software developer for 20 years my take on the matter is simple and the speed of solving bugs is a matter of:

- eliminating the red tape (provide direct contact between developers and the people that send in mail, despite the negative effects this has on performance to produce new features). Thing is that nobody other than the developers is going to deliver on a solution unless the code is made open source. Let the developers share (financially) in resolving the problems, i.e. a developer could get a bonus in a certain promile of revenues of the product in question for resolving an issue. This is similar to assessor fees that are paid by insurance companies or in court cases to assess damages.

- compartmentalize the code as much as possible (somebody mentioned windows and it is too difficult to check within 60 days whether a update is going to break functionality. I believe windows lacks compartmentalization and that this is the cause of things not being fixable)

- consider security a higher priority than functionality. Just like in our society driver licenses are revoked after too many DWI's, we should realize that mobile phone service can seize to operate if that poses a physical harm to people (via whatever way possible, I cannot think of one now). Similarly, flying over Europe seized after the eruption of an Icelandic vulcano. That was security. Though it turned out later this was more strict than required, it was THE SAFE THING TO DO. This namely ties back into the compartmentalization: when features are made available in chunks instead of in one whole block, the functions can be seized and released in managable ways. Service does not need to stop completely for security to be maintained.

- consider that the more our societies are becoming dependent on IT, the more likely it will become that actual damages will be incurred by people and that software liability will become - at some point in time - a consumer right. Loss of functionality no longer will be the only damages that can be claimed, but will include other losses. Though this may actually be very detrimental to the speed of development of software, unfortunately, I see it as an unavoidable outcome given the shameless lack of consumer support by companies producing consumer goods. Let's be honest, I know of no-one that wrote to Microsoft to resolve an security problem at home with a laptop and actually got them to help.
This is why third parties managed to completely dominate that market!

- So far, industrial self-regulation, as the 60 day period you propose, seems only have worked for the smaller companies where there is still direct contact between end-users and developers... Just last weekend I spoke to someone at SAP who was frustrated in having to convince his manager that speaking to the end users was the only way of truely evaluating the companies products. There are IT companies which instantiated proper bug fixing protocols because it was stipulated in the contracts with their clients. For example in the commercial HPC community it is normal that problem response time is < 1 day. Larger fixes need to be available in < 1 week time. I would not be surprised if some EU or US government bodies will impose such regulations also on actual software functionality and software security fixes. After all, if one goes to the dealer and buys an auto, one must be able to expect it to also be able to drive. If such functionality is not present, the car dealer as well as the car manufacturer can be fined for false advertisement.

July 22, 2010 at 8:29 AM

Rogelio said...

For windows it might be better to delay the disclosure. Full disclosure is only going to irritate people. The black hats will find out about vulnerabilities anyway. The main question is, can

microsoft mobilize enough resources to fix the vulnerability in a timely manner? I dont think so. Besides its so buggy anyway that even if they fix it there are a lot more to find right? So whats the use? And its not just the OS that needs to be fixed. Its also the apps. And when you fix bugs on windows, apps that relied on this bugs suddenly stop working. For Linux and the BSD's definitely full disclosure.

July 22, 2010 at 11:09 AM

dremeda said...

How does the responsible disclosure 60 day time frame relate to website based vulnerabilities? Or even more importantly vulnerabilities that have been exploited and are spreading?

I don't think a 60 day time frame is reasonable in the case where a hosting provider has vulnerabilities and major spam or malware outbreak is growing.

Thoughts, comments?

July 22, 2010 at 1:32 PM

William said...

Vendors are taking too long and while they do that the researcher has an effective backdoor in their system. It is great vendors like Microsoft trust security researchers to not abuse or sell or otherwise disclose their zero day while they wait for them to fix it...

Sarcasm.

It is also presumptuous of such vendors to just assume the zero day was not, is not, and will not be discovered and used in the wild by someone else.

(China? Russia? US? Zambia? Whatever. Criminal organizations? Lone wolfs wanting to reboot the world?)

The assumption that zero day might be used and immediately detected is presumptuous.

Zero day can be used for years without detection if deployed carefully by an attacker.

A vulnerability that can effect tens or hundreds of millions of users is a big deal. Especially one that allows a malicious attacker remote, invisible, immediate, and full compromise of these user's systems.

A lot of vulnerabilities are of just this nature which vendors wait on like there is no problem.

People complain about vendor backdoors, but when a researcher has a vulnerability of that critical nature... that is an open backdoor.

There is also the huge need to try and contain the burgeoning black market for such unlocked backdoors. You can not just shove these guys underground because they can make a lot of money with these attacks. They need to be encouraged to go the legitimate route, constantly.

July 22, 2010 at 2:26 PM

Anonymous said...

Tavis Ormandy case was the first new age of cyber terrorism

http://n3td3v.blogspot.com/search/label/Tavis%20Ormandy

Andrew

Security consultant and industry expert

Founder of n3td3v Security

July 22, 2010 at 3:02 PM

Steve Riley said...

*n3tdev says: "Cyber terrorism isn't about killing *people*. It's about cyber attacks on against computers."*

Terrorism's aim is to, uh, terrorize -- that is, to make afraid. When people are the targets, using the word "terrorism" makes sense. "Cyber-terrorism" would be the act of trying to make computers afraid!

What we're talking about here are attacks, pure and simple. And that's how we need to be structuring response processes. Thinking about it as terrorism or as war is just wrong.

July 22, 2010 at 3:04 PM

Matthew Lye said...

Terrorism implies an attack on a people or government.

If you can explain how his disclosure can been seen as a direct attack on a country ill listen. But its not, is merely supplying information that could be used by anyone.

Saying that Full Disclosure is Cyber-Terrorism is like saying that the Anarchists Cookbook is terrorism.

If someone then took the information and used to attack a government, then that person would be a cyber-terrorist.

The improper use of terms just serves to cloud the issue and dilute the meaning of the term.

Also... n3dt3v you really need to go back to school and learn some basic critical thinking skills.

July 25, 2010 at 5:57 PM

Anonymous said...

When we alerted Google of security issues in the past they sometimes didn't consider it an issue -- until the full thing was blogged about (for instance, only after revealed Larry Page's Picasa album which was meant to be private, the wheels at Google got in motion, while before I was being told that such "unlisted" albums could be named cryptically by the user if the user wanted to protect themselves!). Other than that, if no reply comes back at all -- which also happened in the past when I wrote to security@google.com (except for the auto-reply) -- I typically consider 30 days to be fine before a full disclosure (more time may be given if the company gets back and asks for more time).

July 30, 2010 at 9:46 AM

ben said...

Well I recently had my gmail compromised. Not sure if they were using a password cracker or if they were using the asq to get access to the account. I'd almost bet these guys are using an alternative unrated vector into cracking / reseting the asq questions.

I've seen it for sale at http://screen.name/forum/index.php?showforum=69

http://i.imgur.com/8RQ79.jpg
(screenshot since it is a locked forum)

I guess I'll have to buy it from the guy who stole it because google has no interest in restoring my account for me even though I've left my phone number numerous times on the adsense account recovery page.

October 21, 2010 at 6:37 PM

corrector said...

"If it controls a power plant or an MRI system,"

...then the law probably forbids you from sleeping on a critical vulnerability for 60 days. (If it does not, change the law.)

The "testing" argument is idiotic. Obviously, testing was not sufficient to detect the problem in the first place. Only people with no computer science knowledge would think that an OS change has to be tested on every piece of hardware in existence : most OS code either do not deal with hardware at all (like the file system layer, the network layer...), or deal with a specific piece of hardware. Additionally, you CANNOT test every possible hardware combinations. And problems often arise from strange interactions between components that you will not detect unless you can test an insane number of configuration.

I repeat : no amount of testing would be sufficient to guaranty that a patch would not break something. Not in 60 days, not 60 month, not in 60 years!

If things regularly break in some remote component whenever you make a change to some other component, it means your system is crumbling, and that's unfixable. You need to redo it entirely, with more competent designers.

If you need for than 6 days of testing for a simple bug fix, than your product has a problem. (On the other hand, people are already using your crappy OS. Just keep it as crappy so people don't get confused.)

September 7, 2011 at 8:09 PM

Kristian J. said...

Fixing an XSS bug is typically pretty easy as is testing the fix. 60 days seems quite generous for that. Cheers.

November 30, 2011 at 6:20 PM

Unknown said...

I think the ONLY way to keep Your property safe is to backup it often on another machine/server/cloud. We're a victim of huge hacking action, and data loss...

May 14, 2013 at 10:29 AM

batborg said...

I have to say something about Windows all OS , Apple 2+, e, 3, Mac etc, BBC, Trx-80, x OS systems. My lord they are gapping with intrusion flaeed cracks and a noghtmare to control anything. Only One i knew never took a hit but could identify a wannabe for tag and trace. One other thats a beast of a hand held device Motogodlike.* ve been out of check on comp technology for longer than you would believe but I was once a modetator of tech directives, unofficially of course. But from a flip to a touch screen interfaced phone in 5 months I saw a camera phone turn into a tool of power that scared me. Google gave me answers and YouTube music.
After that os a story I wouldn't believe.
Anyway 4 years of technology immersion and now I speak technobabble to no one but the few. Google I trust. The world we're about to walk into is best be experienced rather than spoken from someone like me unless that is what you seek.
This is a time of a new age. How you do is always your choice. How its done will be different. Choice and change. Choose whats right change for you and me a world forall to be free to live and love.
4 OS android phones, 2 OS iphones,
Germany company but location is everyehere. Eyes that see all in detail that I wont describe. Coverage unknown but fair to say global with multi-visual simultanrous imsgery of anything a human with a device with sensors and two eay frequemcy that can reach another device proximity. Signal bounce until water when a satellite or radio wave is required. This is no prob if radio is used but satellites can be digiscanned.
My gyto and both cams were used to a dimensional topography and my movement patterns with all the daya that could tell or fo things easier gor me. Assistant of knowledge help. This is true. Motion, gyroscope, maps, aquasition of interest and listened when I was wondering if anyone was.
They were! They were.
I hope to shake the hand of my friends from a single day years ago.
Guard is always on. Safe we will be.
I know of 2 security giants. If one doesn't know the other then interest will be a part question and database and technology retrived and computational sorting will be a handsome find. And I will request access to all related data. But this is not a black hat. If anything theres stripes involved. So what is write? Hmmmm today's date. Lol.

August 2, 2013 at 2:38 PM

Atmywittsendthanks said...

This hacking and cyber bullying and cyber stalking has been a more than serious issue in my life for over six months!ive talked to my phone carrier my phone had died six times as i have tried to solve this matter alone learning the tech jargon and capabilities as i go i may have inadvertantly caused some of those device failings since i m blindly downloading apps trying to investigate the matter i even put a craigs lis add for a tech assistance only to get contacted twice by two so called techs who i suspect were actually taunting me about what they were actually doing i have flung accusations at friends and family and yea even my boyfrien of putting spy apps that are incogneto on my phone all have denied it and got tired of my accusations continuing since the harrassment has not stopped i am really tired of it all and thanks to all this find myself trusting nobody and suspecting everyone google didnt tell me not one thing about weather or not any of my accounts had settings i needed to change in order to fix any of it. I even pondered permanent solutions to what i hope is a temporary problem but unfortunately i dont see an endto it at all i cant have a phone or anemail accout it seems THANKS EVIL HACKER YOU MUST WORK FOR THE NEW WORL ORDER OR SOMETHING!

June 17, 2014 at 11:41 PM

Post a Comment