

PWNED

THE COLLECTED BLOG POSTS OF TROY HUNT

by Troy Hunt

edited by Rob Conery

© 2022 Troy Hunt & Big Machine, Inc. All rights reserved.

All rights reserved. No part of this publication may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, recording, or other electronic or mechanical methods, without the prior written permission of the publisher, except in the case of brief quotations embodied in critical reviews and certain other noncommercial uses permitted by copyright law. For permission requests, write to the publisher, addressed "Attention: Rob," at the address below.

Published by:

Big Machine, Inc

https://bigmachine.io

Please forward all questions RE publication to rob@bigmachine.io

Version 0.3

i

YOUR FEEDBACK

There is a lot of text in this book and we've done our best to get it right - but as I'm sure you're aware, it's nearly impossible to escape the typo demon! To that end, we've created a place where you can let us know if you find anything. Before we get to that, however, there are some things to know.

The first is that the majority of the text in this book is from <u>Troy's blog</u>. Being a blog and all, we decided to keep it intact, word for word. *There will be typos*, especially in the comments. What we're most keen on getting as polished as possible are the intros and epilogues. Troy just wrote those - so they're fair game.

Next: Troy's Australian and they have a different way of spelling words than their American friends to the North. We used an Australian dictionary for this document so a word that you might think is incorrect could, in fact, be correct.

Troy quotes many sources in his posts. Sometimes those sources spell things wrong - in fact *they often do*. Given that it's a quote we won't be fixing the spelling there.

Finally: when it comes to font choice, layout, legibility, contrast and more - these are all functions of the reader applications that you're reading this text on. Rob's done his best to lay this out as slick as possible - but all of it is overridable by you and your reader.

If you've gotten to this part and you still have something for us, please head over to our <u>GitHub repository</u> and make an issue for us! A PDF page number would be extremely useful so we can find the text you're referring to!

Thanks so much!

EDITOR'S NOTE

My goal with this book is to pry back the professional veneer of "Troy Hunt, International Speaker, Jet Skier and Poster of Amazing Photos" and dig into what really goes on in his life. As public as Troy is regarding the "Good Parts" of his existence, he's also strikingly candid about the other... not so fun stuff. I thought this would be straightforward! It wasn't: I couldn't find a veneer to pry away so... there goes that plan.

I've known Troy for years and have known "of" him (as we say in America) for longer than that. Over the last four years, however, we've bonded over very unfortunate life circumstances.

In June of 2018 I pitched the idea of this book to Troy and he liked it, but he wasn't too sure if people would want to buy something they could read online easily. I then gave him the single sentence pitch that lit him up: *I want to make a book like a This Developer's Life episode where you add context to each post* (This Developer's Life is a podcast I do with my friend Scott Hanselman that deals with the more personal side of the tech world). He liked the idea and we jumped in and got to work and in the months that followed we'd talk from time to time about where we wanted the book to go.

And then Troy disappeared. That was fine with me, I knew he was busy and, if I'm honest, so was I. These things take time.

I let a year go by and in August of 2019 my personal life ... well there's no other way to say it: *exploded*. Maybe *imploded* is a better word for what happened – I really don't know because it was a disaster and possibly both things happened at the same time.

I have never discussed this publicly. I haven't been able to. It's painful, disorienting, and (being honest) I'm still in denial about a lot of it and I have no

ability to handle anyone's thoughts on my personal life, even (I might say *especially*) the nice, well-meaning thoughts from good people letting me know I will make it through all of this and that *one day I'd look back and*...

These are kind thoughts and I do appreciate where people are coming from but, at the same time, I never realized how a kind thought can ruin your day.

Right, this is Troy's book and I run the risk of hijacking it and making it about me. This is not my intent, but I *do* feel you need to understand where I'm coming from as the editor of this book: I care about this damned thing.

In August of 2019 my marriage failed in spectacular fashion, COVID was spinning up and the US political landscape was... well the US political landscape. I was 52 years old, emotionally cornered, alone and *scared*. That was when I decided to go to Australia.

I used to read Alexander and the Terrible, Horrible, No Good, Very Bad Day to my youngest before bed and I guess the message got through to me: if you're facing a life-altering crisis that you're absolutely not equipped for... go to Australia.

So I called Troy. Or, more precisely, I sent him a DM on Twitter:

Think I might need an emergency trip to Gold Coast. My life is turning upside down. When u back? No crying on your shoulder, I promise mate.

Well... maybe just once or twice but after that...

I needed a fast Jet Ski ride and a change of scenery. Maybe we could go wakeboarding on his boat or I could build some Lego things with his kids. I could float in his pool, staring at the blue sky... I was falling apart and I needed to just... float in someone else's pool for a bit.

Troy's reply was unexpected. I don't like sharing DMs unless they're mine, but he said it was OK so here it is:

Uh... my life is turning upside down too. You go first... what's going on?

This was a shock. Troy's life... turning upside down too? How is this possible?

I figured I would find out soon enough so I took a few minutes to assemble the shards of my personal drama into a Twitter DM and let him know what was going on.

My marriage of 20 years had just ended. Therapy failed, *everything* failed. My oldest was going to live with me, my youngest (whom I am extremely close to) was going to live with mom, thousands and thousands of miles away. Our family was being torn in two and there are a few more nasty bits that I'll keep to myself but it really was not a pleasant situation.

Troy's response was swift. I got a series of DMs and a little while later a phone call. I talked, he listened. He talked, I listened. I had no idea anyone could relate to what I was going through, let alone Troy, but indeed... *he could* and he *did*. That call meant everything to me.

A month or so later I flew down to Sydney (for NDC Sydney) and Troy and Charlotte took me out and, like the true friends they are, listened and consoled. Their outpouring of generosity still blows my mind. Even to this day - we meet once a week to talk about this book and Charlotte always makes a point of asking how I'm doing (and my kids too). I do the same for them and we *make time* to check in and see how the other is doing.

Troy gets a lot of flak from various corners of the internet because he has no problem sharing the things he loves about the life he loves. If you follow him on Twitter, you'll notice every third or fourth post is a gorgeous shot of the Gold Coast, his boat, his jet ski, walking on the beach at 6am, a beer in front of his rad desk setup, etc. Many have viewed this kind of thing as excessive, a "look at my fantastic life and despair" kind of thing and, as an American, I can see why people think that. We're a competitive bunch up here in the US and find envy an easy reaction to anyone else's good time. I used to feel the same when I saw Troy's posts – what's this guy trying to prove, anyway?

At this point, in this type of story, you usually read "and then I got to know him", which is true, but it's a lot more than that. Troy and Charlotte took me in

when I needed it most. I have friends that I've known all my life - many of whom I consider "best" friends - that have obvious difficulty when I bring up the things that have happened to me over the last few years. They know my ex, know me, and let's just say it gets weird fast.

I know it will be hard for some people to grasp this, but when Troy shares his amazing pictures, he's sharing his joy and gratefulness *for that moment life has given him*. I don't think the man is capable of strutting around like a tall poppy! At least - I can't reconcile Tall Poppy Troy with Friend Troy who sat across from me at the pub in Sydney, listening quietly as I told my story and then offering me his own in return. Those two people just can't exist in the same mind.

I tell you all of this because I want you to know the level of care that I put into editing this book, which I view as a collection of wonderful stories from my dear friend, Troy. I've read these posts (and their intros) dozens of times - making sure that formatting and image alignment are *just so*. There's so much more to Troy's story than beautiful travel pics and tan legs sprawled out next to his pool.

That said, formatting an ebook is rough business. I've published a few of these things and I can tell you from experience: *they're never finished*. So, if you see something that looks weird, doesn't quite lay out correctly, or is spelled wrong: *it's my fault*.

I Would Love Your Help

This book will be going through a few revisions and if you find a typo, poorly constructed sentence or something that doesn't make sense – for my peace of mind – please let me know! I will happily fix - I want this work to *shine*.

As for me, it's been just over 3 years and the family is adjusting to life separated. I see my kids often (which is great) and my ex and I are on good terms, which is as good as I could hope for. This will take years, but I can feel the scars healing,

slowly. Editing this book along with Troy's and Charlotte's friendship has been a large part of that healing process.

Heavy stuff - but that's life isn't it? There's always more going on than what you might learn from a first encounter or, worse, *Twitter*. Either way: **thank you so much for buying this book**! It's a labor of love and I do hope you enjoy it. There is a lot of life in the following pages.

Rob Conery

June, 2022

FOREWORD: LARS KLINT

On 30 November 2017, Troy Hunt testified in front of the House Committee on Energy and Commerce of the United States Senate. It was a testimony on how data breaches are severely affecting the use and security of any authentication systems on the Internet. I watched it live on my phone lying on my couch at 2 a.m. There was no way in the world that I would miss it. He even wore a suit, which was a necessary departure from the usual shorts and singlet. The testimony was one of the absolute highlights on a whirlwind journey for my good mate, the hardest working person I know.

I first met Troy at a Pluralsight author breakfast during Tech Ed Australia in 2013. I was a newly branded Pluralsight author, yet to release my first course. Troy had released his first two courses and was in the thick of planning the third. I had heard of this Australian security expert and had seen him a few months earlier present at Web Directions in Melbourne. But I was yet to introduce myself in the typical Aussie way: "Owyagoinalright?" That happened at an informal buffet breakfast at the somewhat corny Jupiter Casino on the Gold Coast in Australia. We exchanged small talk pleasantries (something I now know isn't his strong side), but it wasn't until 5 months later that I would come to think of Troy as one of my closest friends.

In February 2014 we both travelled to Salt Lake City in Utah to attend the annual Pluralsight Author Summit. Beforehand Troy had put out a call to the Pluralsight authors to see if anyone wanted to explore winter in the Rocky Mountains for a day before the summit. I was happy to join him and we met up on a cold Monday morning in the lobby of our hotel in downtown Salt Lake City. I won't go to the stretch of calling it love at first sight, but we had an easy time finding topics of discussion and in particular three areas of interest were in common. Technology and how it can make the world a much better place. Cars,

and in particular very fast, competent and exciting ones (of which we have quite some disagreement of the exact definition, once in a while). And lastly, an absolute pathological desire to avoid mindless small talk and "shooting the breeze". This has become, in my mind at least, a defining feature of our friendship. We may only talk briefly once every few weeks, but the content is equivalent of several hours of "normal" conversation. If we discuss a car's performance or appeal, it is because one of us genuinely is considering acquiring one. If we discuss a project's relevance and technology, it is because we value each other's opinion and need a way forward with the tech. If we plan a family holiday, we are keen to hear about pointers about the places we are going. My point is that *all* conversations have a real outcome and it is immensely satisfying.

This book is about the journey of one of the best known internet security people on the planet today, and it starts when security wasn't even a core component of Troy's skillset (and he hadn't broken any jet skis yet). As you will soon discover, the first blog post ever was about establishing a digital presence and footprint for future Troy to use as a base for the other projects, known and unknown, that were to come. Again, there was no small talk, there was no "this is a good way to test my new keyboard", but rather a measured step to become what most people now see.

One of my favourite quotes is "do today what others won't, so tomorrow you can do what others can't" (*Jerry Rice, American Football Hall of Fame receiver*). In other words, put in the hard yards and do the best you can to the best of your abilities and down the track two things will happen. First, many more opportunities will present themselves, and, second, you suddenly find yourself with the knowledge, time and means to do much greater things. As both a spectator to the past years of Troy's journey, as well as a participant at times, this rings very true. There are no coincidences in what you are about to read. Each selected blog post is a good example of a specific point in time, as well as a greater plan.

In April of 2015, Troy left a corporate career in a very large multinational company. Which company and the reasons why he left, you will read about in due course, but the time leading up to that was tense. He had wanted to "do his own thing" for some time, and the year before I had gone 100% freelance, which meant we had some grounds for comparing notes. And there were months of comparing notes. By now it should be clear that the good Mr Hunt will not, and does not, do anything by impulse or in a rush. I admire this, as I am often too impatient to follow through with the big picture plan. In the words of Winston Churchill – "You will never reach your destination if you stop and throw stones at every dog that barks."

With the possession of a schedule that was free from corporate duties, came suddenly a much greater exposure to the public eye. And with the added attention came the trolls and the haters. While people with common sense know that responding to antagonistic behavior online only exacerbates the issue, it doesn't mean abuse and personal attacks has had no impact. Also included in this book is how Troy in general deals with trolls and idiots, but what is often forgotten is that we are all just human. And that includes the good Lord of the Jet Ski. Especially after going independent, the abuse seemed to increase (likely due to people realizing he had worked a full time job while creating more content and sharing more knowledge than they could do in a lifetime), and we had several conversations about why, as well as how to handle this, and it showed me a much more contemplative and demure side. As much as I know Troy wants to control as much of the public image of him as possible, to make sure that people are treated with respect, and facts are put before emotions, sometimes this is difficult. He has shown me plenty of tweets, personal messages and emails of what the abuse is like when no one is looking. I suspect the vast majority of people don't consider this side, nor are even aware that it exists.

With the success of especially the Pluralsight content and "Have I Been Pwned", I saw the pieces of the puzzle coming together. When you pursue a goal of creating content that is as good as it can be, in order to help as many people as possible, the sudden windfall, both financial and professional, comes with a ton

of opportunity. However, at times, many of these can seem overwhelming. The transition to "the next level" is not instant and most people will have a measure of imposter syndrome. While I am sure there were moments Troy would sit on the couch and go "Woah! That was awesome!" about one or another achievement, he would use it as a steppingstone for the next thing. Use the success of the HIBP service to build an API around it. Use the popularity of the Pluralsight courses to do in-person workshops. And still, at no point would he take an arrogant stance, splurge on unnecessary things, or forget the journey and ethics that brought him there. And that is one of the key elements I appreciate in this guy as a mate. He reflects on the many peaks he has climbed and challenges he has conquered, yet still remains grounded.

One of the parts of Troy's life you, dear reader, has most likely not been shown, is the immense stress that comes from walking in Troy's shoes. We all experience stress at times, of course. However, when you are responsible for millions of people's passwords and personal data, constantly in demand for government advice, speaking engagements and corporate engagements, go through an incredibly demanding and stressful acquisition process with Have I Been Pwned, and then add a marriage breakup on top, there isn't much brain power left. Yet, he still "fills in all the rings" on his watch's exercise tracker every day, and he always picks up the phone. Don't know how he does it to be quite honest.

This book is Troy, and Troy is also a very personal and private person. What you might see on social media, what he shares in blog posts, and what he tells about at conferences, is only part of the picture. While Troy is going to give his own account of much of what you don't normally see, in this book, I am going to say this. As much as I admire the guy and his long list of accomplishments, I have no illusions that Troy is easy to live with. I don't know any people personally that work as many hours, and as hard, as Troy does. The content and products he creates takes hours, days and weeks of work, and that time has to come from somewhere. His two kids and fiancée Charlotte are some very patient people. While 2020 and 2021 have seen us all spend much more time at home, his

travelling schedule has otherwise been intense. And speaking from experience, travelling for events, workshops, meetings and more isn't a holiday. You are constantly on, and the days end up being longer than the hours of sunlight in summer in Oslo. We have both written about these trips to give an insight into the massive amount of work they are. On top of the travel, there are the numerous blog posts, videos, podcast appearances, CNN interviews and much more. I know Troy sacrifices a lot of personal time in periods that are particularly crazy. That is a hard combination in any family.

Despite the many hours that Troy puts in each week, and despite the much greater demand on his time in the past few years, he always has time for a quick chat or to discuss an idea either of us might have. The Troy I met in 2013 is still the same today. He doesn't do small talk. He gives you an honest opinion. He doesn't compromise on ethics and quality in the projects he does. And he still rolls his eyes at my amazing jokes. I will get him to appreciate them soon though.

In a sense I am a bit jealous about the experience of reading this journey that you are about to have. You get a curated and hand-picked view of how Troy Hunt went from enterprise IT employee to international web security expert testifying in front of the United States Congress. I have been part of most of that journey, but often the view has been cluttered or hazed with my own journey, or the dead ends we all invariably pursue. You get the best bits, and the most valuable parts of what has made Troy able to do what he does today. And the best part? There is so much more to come. Troy is never short of grand ideas or projects. Just don't give him another jet ski.

FOREWORD: RICHARD CAMPBELL

Look, being a security person isn't easy. Troy may make it look easy, but it isn't. It's a constant battle of people not taking you seriously, suffering the consequences, and saying "I told you so" doesn't help.

I remember consulting with one organization that called their InfoSec guy "the business impediment service." The problem is, when you spend all your time focused on security, actually starting to understand how serious the situation is today, well, it's hard not to get a bit paranoid. Suddenly, the tinfoil hat doesn't seem so outrageous.

And when you see so many mistakes happening in security, so many breaches, ransomware attacks, and so on... it's easy to become a bit self-righteous. The "I told you so" is never far away.

That's what I noticed about Troy when I met him: The total lack of self-righteousness.

The first time Troy was on .NET Rocks was back in 2012 – before Have I Been Pwned, or Pluralsight, or any of the craziness today. I remember teasing him about scaring the listeners with all his security concerns. He talked seriously, head up, eyes open, and truthfully. But, in hindsight, I think he underplayed the reality.

When HIBP first went live, I signed up – and immediately received notification that I was on a list. Unfortunately, I was part of the Stratfor hack back in December of 2011. At the time, breaches were rare, and I changed my password and carried on, not thinking much about it. Stratfor also offered anyone caught up in the breach a year's worth of identity monitoring, which seemed generous at the time.

When the next email arrived about the Adobe breach, I decided to take my

security more seriously and signed up for a password manager. It doesn't matter which one: Any password manager is better than no password manager.

All these years later, I realize now how much stress it has removed from my life. I don't know my passwords – they're all different, and they change regularly. I love how my password manager helps me track what websites and applications I'm using and what ones I haven't touched for years. It feels good to delete old identities. Do they actually delete them? I don't know.

Listen - you want a password manager. Your life will be better. It will take a few weeks to get a feel for it, and it will take a year to get all your passwords different. But you'll be happier. If there's one thing that HIBP has taught me, it's to use a password manager. It's not as if data theft is going out of style.

HIBP was sort of a lark in those early days, really – a reminder that hacking was real, data was being stolen every day. And somehow Troy ended up in this inbetween space, between the exploiters and the exploited. We, as the exploited, counted on HIBP to let us know when our information had been stolen (again). And amazingly, the exploiters also depended on Troy. Getting your breach listed on HIBP was a kind of badge of honor – it was proof that you had done a thing, as malicious as that thing was.

All of that changed with the Ashley Madison hack.

The concept of the site is grotesque to me but to each their own. Remarkably, the hack revealed that the site was essentially a lie and that while life may be short, it doesn't look like anyone's having an affair, at least through Ashley Madison. Lovely fodder for the gossip pages and an easy opportunity for anyone to look down on the immortality of others. It astonishes me that after the breach and reveal, the site continues to exist.

But as disingenuous as Ashley Madison is, the impact of the breach was all too real – this was the hack that destroyed lives.

I remember spending a day with Troy when the Ashley Madison breach was surfacing – and the pain he was in from the suffering he was witnessing. The endless emails from people begging Troy to take them off the list, accusations from spouses – even churches using the data to check up on their parishioners. He did his best to help those that reached out to him, but it was overwhelming. HIBP may have been the key to Troy becoming a titan – but it had also become a weight upon his shoulders he could not put down.

Troy's solution seems obvious today, but that doesn't mean it was easy – it took a lot of thought and careful consideration to help protect people with the truth rather than harm them with it.

But the impact that Troy was having on the understanding of the security crisis was only starting to show – testifying to the US Congress (you know he's an Australian, right?), chatting with Philly DeFranco on YouTube. Our Troy gets around. And his message is remarkably consistent – we need to do a better job of protecting our own, and our customer's, information.

Troy has also taught us the right way to deal with a breach – head up, eyes open, speaking truthfully.

I feel only empathy today when I see a tweet from Troy along the lines of "Does anyone have a contact at XYZ company?"

Let's face it. It's never good news. But as bad as the news might be, it's not going away if you ignore it.

Neither is this security crisis. HIBP has only informed us how serious the state of cybersecurity is. It's up to all of us to deal with it – head up, eyes open, and speaking truthfully. It's the least we can do, and it's the only way to get beyond the crisis we're in now.

It also helps to look how far we've come - and in some ways, Troy's story is our story - of taking the problem seriously and finding ways to convince others to do so. So read on!

INTRODUCTION

magine me when I was 18 on the cover of this book. It was 1995, I'd just seen the internet for the first time and suddenly, there were all these possibilities I'd never dreamed of before. There was this sense of... empowerment. The idea that I could sit at home on a PC and write code that anyone anywhere in the world could see without needing to ship it on a CD! Mind. Blown. Without knowing it at the time, the next few years would define the next few decades and make me rich, then poor and ultimately set me on the path I find myself on today.

I'd started university that year and the idea was that's what you do to pursue a professional career. Computer science seemed like the logical path to follow, and my mediocre high school grades were just enough to get me into a uni close enough to the beach that I could split my time between doing what society expected me to do and what I actually wanted to do. For the most part, I disliked formal education for both pragmatic and philosophical reasons; I couldn't find courses on the topics I actually wanted to learn and if I'm honest, I didn't like people telling me what to do then assessing me on it. I wanted to feed that emerging excitement I had for the world wide web but there was no chance of finding content on anything even remotely related. Discrete mathematics? Yep. Chemistry? Ugh, and yep. So, I either scraped through or in that case of the latter, failed.

HTML for Dummies for my saviour. I mean literally the classic yellow book series which these days, will teach you everything from how to play poker to how to grow pot. In '95, that was my best bet because everything about the web was so new. But that's what made it so exciting because it was new to everyone, and we were all just hacking away at code trying to do things that had never been done before. Initially that was just out of curiosity, but the penny soon

dropped that *people would actually pay me money to build web pages!* Cool. A car hire website. A brochureware site for a stable. And then... the horse racing software.

It began with Adrian. He was a bookmaker (someone who takes bets on sporting events) and horses were his thing. He knew enough about the ins and outs of the races that he'd developed a formula that *guaranteed* he could make money, and he wanted to share that with the world (for an upfront "investment" on their behalf, of course). But as much as he knew about nags, he knew nothing about code and in order to fulfil his vision of sharing horse racing prosperity with the masses, he needed someone to write the software for him. But that wasn't me – not solely – because HTML for Dummies and the vague attention I paid to my courses wasn't going to cut it.

Angela paid a lot more attention in class than I did. By now we were well into '96 and I was in a relationship with a smart, beautiful woman who could code. It turned out that some of those uni lessons actually contained some useful information and she'd listened closely enough to be able to write the back end of Adrian's vision in C, committing the data to flat files for the transactions. I complimented her work by building the front end in HTML and JavaScript (there was no CSS yet and HTML 4 was still a year away). We stayed up late at each other's' parents' houses, cutting code we "versioned" by copying folders and naming them based on the date. Eventually, we had a minimum viable product we delivered to Adrian, and that, with the benefit of hindsight, is when our fate was sealed.

He brought on John to sell the thing. Apparently, he was the guy that could take Adrian's formula plus Angela's and my code and turn it into a marketable product. Not just a product, but an *investment* that could be sold to punters who would then share in the aforementioned prosperity. John was the epitome of a salesman; get yourself a mental image of what a self-confident, brash, pushy spruiker looks like and that was John. As you'd expect from a character of this type, he promised the world and we began having visions of success that were previously unimaginable for a couple of near-kids at university. But I was also

hesitant and frankly, I didn't completely trust either of the older guys not to take advantage of Angela and me. And so, to protect our interests (and this is where I made an absolutely horrible mistake), I insisted on being a director of the company. A director has control, right? And power – the ability to stop other people in the company overriding us, correct? Well, kinda, in a way, but a director also has *accountability* and that's something I really didn't grasp at the time.

We moved forward with us coding, John selling and Adrian placing the bets. It started very small in a tiny, serviced office, then a larger space and before we knew it, two full floors we'd bought in a building nearby. Adrian had a team of people helping him with the betting. John had what felt like an entire floor of salespeople. Angela and I had half a dozen techies. And, of course, we were making money. Lots of money. I moved out of home and rented an apartment, decking it out with new furniture and hifi gear. We both bought new cars, albeit on finance. We made A\$800k in a year. A\$800k – in the late 90's – in our early 20's!!! When that much prosperity comes so quickly and so early in life, it's easy to lose a bit of touch with reality, and the reality was that we'd gotten way too deep into something that was about to blow up.

For reasons that to this day I still don't understand that headline figure was just that – a headline – and the *real* figure we received was about a quarter of that. That alone would be bad news, the idea that three quarters of your income had gone elsewhere, but what was much worse again was that we ended up being on the hook for tax on the entire figure. As best I could understand it both then and now, the organisation had been structured such that the other parties were able to grab not just their share but a big chunk of ours as well and have us liable for the tax. But it was much, much worse than that.

Betting on horses is one thing. Betting on them for other people is another. And something else altogether is marketing a product as an "investment" because as it turns out, there are some pretty strict government controls over that and the corporate watchdog down here in Australia wasn't happy about it at all. So, they

opened an investigation. I don't know which came first but around the same time, "investors" were becoming disgruntled and a bunch of them that wanted to pull their money out couldn't get access to it. We'd built software that would track customers, bets, winnings, and payouts, but by design it gave the operators enough free reign to enter whatever figures they liked into the system. Looking back at it now with mature eyes in a post-Madoff era, I can draw some reasonable conclusions on what might have been done with the numbers...

Everything started to come crashing down at once; Angela and I couldn't pay our tax bill, the Australian Securities and Investment Commission was investigating the operation and customers were demanding money. Other creditors were also knocking on the door, everyone from the companies that had leased us equipment to the ones we had car loans with to the banks that had financed the office purchase. And to make matters worse, because my youthful naivety had caused me to be a company director, I was accountable for the company's actions regardless of how informed I was of what was actually going on. I was fucked.

Everything fell apart in quick succession late '98 and into '99. My relationship, my career, my financial position, and my credit record. It nearly bankrupted me and whilst my personal debts weren't too severe, the legal bills I racked up whilst dealing with the whole mess *really* stung. It took the best part of a decade to fully recover, eventually putting all the debts behind me, clearing my credit record and finally being able to look forward and not back.

That story helps explain this story, the one you'll read in this book. So much of me from that time carriers through into the person who wrote the decade and a bit worth of blog posts you'll read here, starting with the wonder I still feel when I sit down and create something for the web. It's exactly the same thing 2 and a half decades later where I still have this sense of amazement that I can sit down here at this very PC, create something for next to nothing then put it out for the world to see. That's how I felt when I built Have I Been Pwned in 2013. It's the same with my relentless tendency to make stuff work (or break stuff!)

where there's this thing I want to achieve, and I just plug away at it until it's done. And finally, you'll see that same drive to push through hard times – *the worst of times* – which 20 years after the horse racing fiasco, I faced again with divorce. Always be looking forward, forge ahead, get over it and move on to happier days as fast as possible.

Enjoy the book 😊.

WHY ONLINE IDENTITIES ARE SMART CARFFR MOVES

I had no idea what I was doing when I wrote this. A mate of mine at work had encouraged me to begin writing a blog as he himself was doing at the time. If a sensible person had looked at my life then, it really didn't make any sense to embark on a commitment like this; I was both happy and busy at my job and, more importantly as it related to my free time, I had my first child due only a couple of weeks later. But I had an itch I wanted to scratch, something that bugged me and that I wanted to articulate in writing. Over the years, that would become the genesis for so many of my blog posts, namely that something was pissing me off and writing about it made me feel better. The trouble I was having with recruiting people was the spark I needed to write my first blog post and then, well, everything kinda escalated from there...

27 SEPTEMBER 2009

he final catalyst for me eventually taking the leap into the blogosphere came from an unexpected source. It was actually my own response to a Stack Overflow Question where I'd suggested that one of the best ways to make yourself more marketable as a software developer is to have an active online profile. I don't necessarily mean to try and achieve semi celebrity status like Scott Guthrie or Joel Spolsky, rather to be able to illustrate that over time, you've been actively involved in the areas in which you profess to have expertise. It's one thing to present a CV or a LinkedIn profile which says you've done everything from writing enterprise software to creating perpetual motion, it's

quite another to be able to reliably substantiate it.

Why is it important to me?

Let me get one thing cleared up right away; I'm not looking to change my job in the near future and for the most part, I enjoy what I do. The thing is though, building an online profile is not an overnight process and I don't know if I'm still going to be as enamored with my job (or my employer as enamored with me!) in two years, five years, ten years; whatever! It takes a lot of time to build a public identity and waiting until you actually need one is just not going to work.

The LinkedIn LoveIn

There are a couple of big issues with the current system when it comes to people marketing themselves. The first is the self-ingratiating, reciprocal backslapping that goes on with <u>LinkedIn</u>, a practice I'm coining the "LinkedIn LoveIn". The LinkedIn LoveIn surfaces itself through the recommendation system where person **A** espouses how fantastic person **B** is in return for an equally impressive recommendation from the recipient. Worse still, the practice is rampant between people who have been unceremoniously shuffled out of their previous employment.



This is all very warm and fuzzy for the participants involved in this practice but it doesn't do a lot for the potential employer who later interviews them, as although the reports are all very glowing, they're very difficult to substantiate. Of course this problem is not new, it's just much easier to propagate now as it only takes a couple of clicks to get a recommendation from someone who can carefully construct a glowing report in their own time as opposed to the old fashioned way of being questioned on the phone and needing to provide answers off the top of their head. And to that effect, when a phone reference is done it's usually with the subject's superior unlike the LinkedIn model which is essentially a free for all.

And just in case you were thinking LinkedIn recommendations are really only a problem for future employers, here's a quote from the National Law Journal in an article titled <u>Lawyers warn employers against giving glowing reviews on LinkedIn</u>:

Plaintiffs' lawyers, they fear, are scouring these sites, looking for evidence to dispute firings, as most LinkedIn recommendations are positive. So if a supervisor claims that an employee was let go due to performance problems but gave a rave review about him or her on LinkedIn -- that, the lawyers stress, won't look so good.

Certainly makes you think twice before going out on the public record and waxing lyrical about someone's performance. And yes, I know, the risk here is more if you're directly involved in someone's departure on the one hand and

slapping them on the back with the other but either way, you want to be pretty damn sure about who you're recommending and what the repercussions can be further down the line.

Setting the bar very, very low

The next big issue is just how low the competence bar seems to have been set in the software industry. The number of times I've interviewed people and they struggled with the most fundamental of questions is staggering. Granted, recruitment agencies have a lot of blame to share but at the end of the day if you're calling yourself a senior .NET developer and can't even write code to declare a nullable type or instantiate a generic collection, you've got issues.

It's not just syntactic ability either, it's general awareness of the industry. I tend to ask a lot of questions about what has changed between versions of technology the interviewee professes to have expertise across or what might be in the future pipeline and very frequently I'm met with a blank stare. In many cases people just don't seem to have an awareness of the concepts many of us take for granted. In short, there's a lot of <u>unconscious incompetence</u> floating around.

The importance of an online identity

It's very hard to consistently fake competence over a long period of time through an online identity, certainly if it involves discussion with a community of peers. That's not to say that every word someone puts out in the public domain should demonstrate their superiority in the subject of the day, it's simply that through their online identity a person discloses a certain amount of information about their competencies. There's nothing wrong with .NET developer asking about how to build their first "Hello World" application in SharePoint and this sort of

active information seeking is great, just don't come looking for a job as a senior MOSS developer the next week!

Building an identity

The way I see it, you've got three key avenues to create an identity for yourself these days:

- Twitter; probably goes without saying given it seems like every second person is tweeting these days, but Twitter is about the easiest way there is to get yourself out there.
- Forums; sites like Stack Overflow are a great way to build up a public profile and of course you're helping your fellow professional as well.
- Blogging; sure, this takes a lot more work than the first two but also gives you a pulpit not limited to 140 characters or a discrete topic.

If I can't find any information about someone whatsoever in any of these sources, it does start to make me wonder. At best the person doesn't tend to use these medium as they're simply passive online users, but at worst, they simply haven't been as active in their professed subject matter as they'd like you to think. Either way, no online identity means you're left wondering.

Recently a friend of mine told me about the interview he went through for his current job. He had difficulty answering some questions on the spot and tying them back to actual experience he'd had. He felt he'd bombed. It was much later on he found out the only reason he got the job was that his employer was able to substantiate, through his blog, that the guy really knew what he was talking about and had simply had a bad interview.

My identity

I'm pretty clear about what it is I think I do well and the general sort of thing I want to be doing in the immediate future (at least technology wise). I've had the better part of a dozen years actively involved in coding and while I've really enjoyed it (and *still* really enjoy it), it's not something I do a lot of any more and quite frankly there are people out there who do it a lot better than me (refer to some of the other blog links on the right hand side). So I'm not going to focus much on actual code in this blog and the syntax I do post will probably be pretty rudimentary.

My online identity will focus on the more practical use of software within business and enabling others to deliver it effectively. I've found myself spending more and more time lately working to try and bridge the gap between how software developers like to work and the expectations enterprise has of them. And that's a two way street; developers are generally not very well understood by those not actively involved in the coding process but by the same token, developers frequently have trouble relating their work back to something that makes real commercial sense.

So that's what I'm setting out to achieve with my identity; bring some sanity to the developer / business relationship, try and show it's something I'm actively involved in and have some idea about and all things going well, not need to rely on LinkedIn recommendations further down the track!

Comments

Hi Troy,

I'm getting my online identity sorted with a professional blog to enable future career

prospects. Your identity is obviously based on your name, what do you think is best, using your name or a pseudonym? And why? i.e. <u>troyhunt.com</u> versus <u>thesecurityexpert.com</u>?

Troy: Because other than that domain name being way too self-ingratiating, you might want to specialise in something different one day. I didn't want to be defined by the thing I'm interested in today.

Ні Тгоу,

thank you for replying so promptly, I'm not sure I'd consider that title too self-ingratiating for yourself but I do agree with the premise that you'd end up switching around the place depending on interests. Thank you for helping cement that conflict in my brain!

Did you ever consider going for the pseudoname approach or were you always fixed on sticking with your name?

Troy: Well it *definitely* would have been too self-ingratiating when I actually started the blog 8 years ago! But regardless, I get your question about branding based on what you do rather than who you are. I never considered this route, I've seen some people build up a brand around a pseudonym but I'd much rather link what I do to my own identity. I think this reinforces who you are and makes you more marketable as an individual.

Epilogue

I love looking back at this blog post now because the theory I proposed in the title became not just my own personal reality, but the thing that helped me achieve dreams I'd had since childhood. From that first blog post, my profile

grew, I gained independence, travelled the world and carved out a livelihood that all began by creating an online identity.

A full 8 years after writing this post, I found myself on the main stage at Pluralsight Live in Utah keynoting next to Scott Guthrie and Joel Spolsky. I reflected on the comment about not trying to "achieve semi celebrity status like Scott Guthrie or Joel Spolsky" and smiled. I loved that I was there and that it'd happened organically rather than being some sort of well thought out plan I'd executed over many years.

What I love most about this post is that the thesis of someone - anyone - being able to sit at home, put stuff out on the internet and go on to do great things held true. I'd had that spark in me since the very early days of the internet when published my first website back in '95. Being on the web felt empowering, like I had this enormous leverage that gave me the potential to achieve so much more than what the resources I had access to would have otherwise allowed. Ultimately, the stuff I put out on the web ranged from blog posts to training material to Have I Been Pwned, all of which were done with little to no resources and all of which began with that single blog post.

THE ONLY SECURE PASSWORD IS THE ONE YOU CAN'T REMEMBER

I remember the first ever podcast I was invited onto. It was the middle of 2010, and it was run by a local Microsoft MVP in Sydney, Richard Banks. The podcast was called "Talking Shop Downunder" and I was there to talk about developers and security. This was really early days for me (we're talking 9 months into my blogging career, with just a handful touching on security), and I distinctly recall being asked about my views on password managers. I wasn't keen – "all the eggs in one basket" sort of views. I feel comfortable admitting that here as it was part of my own personal development to explore different concepts and viewpoints before having sufficient experience to form an educated opinion.

Fast forward to Christmas that year and with a couple of weeks of downtime, I started exploring password managers. It formed part of this gradual transformation I was going through, coming from having a blog and little idea what to do with it, to beginning to focus on infosec in general and very often, specifically passwords. I've always been the type of person who needs to experience things for myself in order to both understand them and get endorsed in them, so I started porting all my passwords into 1Password and never looked back. It was an absolutely pivotal moment in my career with so much of what I later did tracing its origins back to the realisation that the only way humans can consistently have decent passwords is by using a password manager.

21 MARCH 2011

et's assume you log onto a bunch of different websites; Facebook, Gmail, eBay, PayPal probably some banking, maybe a few discussion forums and probably much, much more.

- Do you always create unique passwords such that you never use the same one twice? Ever?
- Do your passwords always use different character types such as uppercase and lowercase letters, numbers and punctuation? Are they "strong"?

If you can't answer "yes" to both these questions, you've got yourself a problem. But the thing is, there is simply no way you can remember all your unique, strong passwords and the sooner you recognise this, the sooner you can embrace a more secure alternative.

Let me help demonstrate the problem; I'll show you what happens when you reuse or create weak passwords based on some real world examples which should really hit home. I'll also show you how to overcome these problems with a good password manager so it's not all bad news, *unless you're trying to remember your passwords*.

The tyranny of multiple accounts

Think about it; how many accounts do you have out there on the internet? 10? 20? 50? I identified 90 of mine recently and there are many more I've simply forgotten about. There is absolutely no way, even with only 10 accounts, you can create passwords that are strong, unique and memorable.

What happens is that people revert to patterns including family names, pets, hobbies and all sorts of natural, somewhat predictable criteria. Patterns are a

double-edged sword in that whilst they're memorable, they also predictable so even if the pattern might seem obscure, once it's known, well, you've got a bit of a problem.

Patterns and predictable words are bad, but what's even worse is password reuse. Because we simply end up with so many of the damn things, the problem of memorising them gets addressed by being repetitive. Easy? Yes. Secure? No way.

The problem with weak passwords

Firstly, what exactly is a weak password? Let me answer this in a roundabout way by focusing on strong passwords; a strong password is one which has a high degree of what we call <u>entropy</u>, or in simple terms, one that is as long and as random (in terms of both character types and sequence), as possible. As the entropy link explains:

People are notoriously remiss at achieving sufficient entropy to produce satisfactory passwords.

People struggle with strong password because they revert to patterns that are easily memorable. The patterns may be in a natural form such as someone's name, a date, or a place or they may be memorable keyboard patterns such as "qwerty" or "123456". These are all highly predictable patterns.

Let me demonstrate the problem with this based on a few recent events. Firstly we have Gawker who last December were the <u>victims of an attack</u> which lead to the disclosure of somewhere in the order of one million user accounts. Worse still, these accounts were posted online and readily accessible by anyone who wanted to take a look at who had signed up to the service and what their password was.

The interesting thing in the context of password strength is the prevalence of

bad password choices. Take a look at these:

123456, password, 12345678, qwerty, abc123, 12345, monkey, 111111, consumer, letmein, 1234, dragon, trustno1, baseball, gizmodo, whatever, superman, 1234567, sunshine, iloveyou, fuckyou, starwars, shadow, princess, cheese

These 25 passwords were <u>used a total of 13,411 times by people with Gawker accounts</u>. The first one – 123456 – was used over two and a half thousand times alone.

Another very similar example was an attack last month on rootkit.com. Password analysis on the breached database showed these top 25 passwords:

123456, password, rootkit, 111111, 12345678, qwerty, 123456789, 123123, qwertyui, letmein, 12345, 1234, abc123, dvcfghyt, 0, r00tk1t, ìînêâà, 1234567, 1234567890, 123, fuckyou, 11111111, master, aaaaaa, 1qaz2wsx

Look familiar? Worse still, you can easily see the corresponding username <u>if you know where to look</u> (I've deliberately blurred these but the originals are still there in the link):

name	email	cleartext_password	
therioox		foofoo	
Edward W. Ray	grampal@mercinan.com	1satriani	
		1234567	
	diplimatins	datalife	
Norths	remark to the after contract	mjlovemo	
ne*freation	natmaniar@hotmat.log	adik1981	
Haurs Faredes	man produdgration	mang1729	
	Section (ECC) com	19790809	
Solomak Salahadis	tobul@seaffbrox.jp	abomb001	
	administration of the community of the c	141421	
Dragon Krome	Bragan, Iromedication com	kokolino	
de tribune	selforms james (Bygmalicom)	guru11	
	Saltes (PRopers com	butterfly	
Dan H	remember handsworth@hushmad.com	bolivar	
Legitor Depart	legition, these or lith street con-	l9s8d78	
to distant	bodnator@cohel.org	cryptman	
Fallence Sharman		bond007	
	confounce # 1015 com	801225	
randh	sand-dhatnal con	ah002100	
Michael Davis	miles (b) dataments met	test	
Roam	Markenshram at , rober rights risk	tubular	
CEPRK	citigation disproals com-	s3gf4ult	
Stephanie (pallechas		insanity	
Chris Wedner	chrises@linelcoud.natl	weber91	

But here's what's really interesting about both these cases and the relevance to why password strength is important – all of these were stored in an encrypted fashion in the database. Without delving into cryptography concepts, the crux of the problem with both these sites is that the encryption was implemented badly.

When a database such as rootkit.com is released into the wild with poorly implemented encryption, hackers are able to recreate the encryption process by feeding in a dictionary of common passwords and attempting to compare them to the database to find matches. The nature of encryption can mean this process needs to be repeated millions of times, but it's an entirely automated process.

Password dictionaries are commonly available (wonder if you see any of yours in

there?), as is the software to run them against the breached database. The biggest limitation is the computing power required to perform a fairly resource intensive process but as we all know, compute power is increasing at a very rapid pace and besides, you can easily acquire enough processing power to test 400,000 passwords per second for only 28 cents per minute.

But the bottom line is this; if your password conforms to a recognisable pattern, there's a good chance it will either be in a password dictionary or guessable based on other known information about you (wife's or kids name, etc.) If it is short or doesn't contain sufficient variations in characters, the number of attempts required to guess it are going to be much lower; you become the <u>low hanging fruit</u>.

The problem with password reuse

You're probably already aware that you shouldn't be reusing the same password in multiple locations, but let me illustrate as clearly as I can, from a firsthand perspective, why not. Here's what was waiting for me in my email when I logged on recently:



Dear Trapster User:

The Trapster team has recently learned that our website has been the target of a hacking attempt, and it is possible that your email address and password were compromised. We have taken, and continue to take, preventative measures to avoid future incidents but we are recommending that you change your Trapster password. As always, Trapster recommends that you use distinctive passwords for each site you visit, but if you use the same password on Trapster that you use on other services, we recommend that you change your password on those services as well.

For information on how to reset your password or improve the security of your passwords for your Internet usage, please click <u>FAQs</u>.

Sincerely,

The Trapster Team

In case it's not perfectly clear, having *your email address and password compromised* isn't exactly ideal. When the scope of those credentials is one website, it's an inconvenience. However, if those credentials were reused across your financial institutions, your social networking sites or particularly your email account, that's not inconvenient, that's downright scary and potentially very expensive for both your hip pocket and your reputation.

Only the day after the Trapster incident, tweets like this started popping up:



Found out my #Google account was hacked b/c of the #Trapster app I downloaded & deleted 2 years ago. Just awesome.



Going back to the Gawker incident I mentioned earlier, shortly afterwards, something odd started happening to the Twitter accounts of people who also had accounts with Gawker; they started <u>ranting on about Acai berries</u>.

This is a crystal clear example of what happens when you reuse credentials. The Gawker database was large enough and the whole password reuse phenomenon rampant enough that the perpetrators *were bound to compromise a lot of Twitter accounts*. What these incidents are showing us is that <u>based on real-world data analysis</u>, <u>password reuse is alarmingly high</u>.

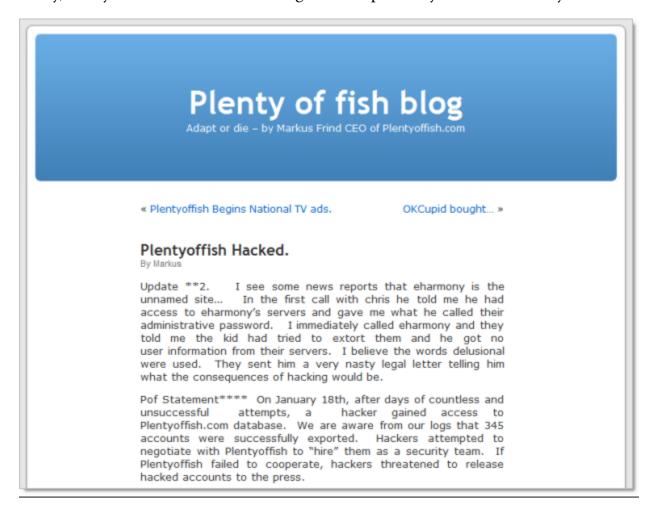
Undoubtedly, much of this problem is related to poor security implementations on websites. It's very, very easy to build websites with fundamental security flaws. Another problem in this area is that all too often software developers take the attitude of "The information on our site isn't that sensitive so security isn't too important". Of course if you've gone and used the same credentials for that site and your PayPal account, you could have a serious problem just around the corner.

Because we all reuse usernames – and often your username is your email address so there's not much choice – it's a very short hop from one compromised account as a result of a database disclosure to another compromised account simply by matching usernames and passwords. In fact there's a school of thought that <u>usernames betray you</u> and <u>Hotmail even recently gave you the ability to easily create additional email addresses</u> which could mitigate the risk of matching accounts but that's probably going a little further than what you

really need to right now.

Just how prevalent is this sort of thing?

Very. Gawker, rootkit.com and Trapster are all very recent examples but there are many, many more. Into online dating? You've probably heard of "Plenty of Fish":



Like the scented, soapy goodness from Lush? Their UK site got hit earlier this year:

Attacks on Lush website expose credit-card details

By Matthew Broersma, ZDNet UK, 21 January, 2011 17:24

Topics

Credit card, Breach, Hacking, Hacker, Ecommerce, Cosmetics, Security, Notification, PayPal, Website

Sponsored Links

Quick and Easy Loans Get Cash Loan or Cash Interest Rate, Call Us

Credit Card Comparisons

& Apply for Credit Cards Online at Mozo today!

NEWS Cosmetics company Lush has warned customers that its UK website has been hacked repeatedly over the past three months, exposing credit-card details to fraudulent use.

Lush did not release technical details of the attack, nor specify the number of customers compromised or the security techniques used to handle the data involved, but anecdotal evidence indicates that some customers have been the victims of fraud.

The company sent an email statement to customers on Thursday outlining the incident Advance in Minutes. Low and urging them to contact their banks.

www.LoanCentres.com.a "Our website has been the victim of hackers," Lush said in the email. "Twenty-four-hour security monitoring has shown us that we are Find Reviews, Compare still being targeted, and there are continuing attempts to re-enter. We refuse to put our



The website of cosmetics retailer Lush has been hacked repeatedly over the last three months.

Photo credit: Kake Pugh on Flickr

customers at risk of another entry - so have decided to completely retire this version of our website."

Not in the UK and think your Lush details are safe? Not quite (but don't worry, the incidents are "unrelated"...):

FRESH HANDMADE COSMETICS

LUSH WEBSITE PRIVACY BREACH Our website has been the target of hackers

We are sorry to have to announce that the Australian and New Zealand websites have been hacked. We were alerted on Monday 14 February 2011 to advise us that entry has been gained and customer personal data may have been obtained by the hackers

We urgently advise customers who have placed an online order with Lush Australia and New Zealand to contact their bank to discuss if cancelling their credit cards is advisable.

Whilst our website is not linked to the Lush UK website, which was recently compromised, it appears that the Australian and New Zealand Lush sites have also been targeted. As a precautionary matter we have removed access to our website while we carry out further security checks.

Lush is working with the police, forensic investigators and banks and doing all that we can to investigate the breach in privacy. We are currently in the process of contacting each of our online customers individually by email.

Of course these were all very targeted attacks. Malicious computer activity goes

well beyond this and is often very indiscriminate. We're now at about <u>50 million</u> viruses and counting, 20 million of those having hit people just last year.

I'm making these points not to scare you, rather I'm trying to make it evident that this is a very, very common thing indeed. The examples above are just a few of the ones we actually know of from very recent times. There's a significant order of magnitude more where your credentials have been exposed that we don't know of, and probably a good proportion of those where the website operators don't even know of the breach. This is commonplace folks, and it's up to you to make a preemptive strike against the bad guys.

The myths of "secure" passwords

First and foremost, the word "secure" is frequently thrown around like it's an absolute term. It's not. Look no further than the <u>Stuxnet virus</u>; computers running the centrifuges in Iranian nuclear facilities entirely disconnected from the internet were successfully targeted by the virus. Surely those systems would have been considered "secure" by any reasonable definition of the word.

It's a little bit like saying a car is "safe". Some are better than others, no doubt, but at the end of the day it becomes a risk mitigation exercise. You trade some things off – such as the simplicity of a password or price paid for a car – and you get a better risk profile in return such as longer to crack the password or more airbags in the car.

<u>Here's how some people (Google, in this case)</u>, believe you should create – and remember – secure passwords: https://youtu.be/0RCsHJfHL 4



Seriously? Can you imagine trying to remember dozens of "I love sandwiches" style of passwords? Keep in mind you need to remember what the phrase was, which characters you substituted and which one you used for which site.

Besides, the whole idea of strong passwords is to avoid predictable patterns. Is substituting an "@" in place of an "a", or a "3" in place of an "e" *really* going to throw the bad guys off the scent? Memorised patterns with substituted characters are a very thin veneer of security and trust me, the bad guys have heard of this trick.

In fact, the password dictionary I linked to earlier contains many common occurrences of character substitution. In there you'll find examples such as "s@yg00dbye" and "s0cc3rRul3s" – not exactly "secure".

Writing your passwords down on paper also isn't going to do you any favours. Because you've got so many of them (and face it, you do), you're going to need to also write down which account the password belongs to which means you've got the mother lode of credentials sitting there ripe for the burglar / kids / nosy guests.

The other problem with handwritten account details is that these days many of us are logging in too many different locations such as the home PC, work PC

and increasingly, our mobile devices. We can't practically have the keys to our online world locked away in a drawer somewhere – it's simply too big of an inconvenience for many people.

And finally, the handwritten *strong* password is just too damn painful to continually re-enter every time you logon somewhere. Remember, a strong password is very long and very random; exactly the attributes which makes manually typing them tedious and error prone.

So, what about just storing them in a Word doc or in a notes system like Outlook? Because they're just too easy to steal and when this happens, they're easy to extract because they're not encrypted. Someone gets their hands on that file and you are well and truly compromised in a most unpleasant way.

Liberating yourself from the tyranny of passwords

At face value the title of this post sounds odd. How on earth can you continue logging on to websites if you've forgotten all your passwords?! You need a dedicated password management system, pure and simple. There is just not another practical and secure way of dealing with it in the current day.

Fortunately, there are tools out there focused at doing just that. For example, there's <u>LastPass</u>, <u>KeePass</u> and my personal favourite, <u>1Password</u>. All of these tools give you the ability to record all your passwords in a single, strongly encrypted location. Of course you still need a password in order to unlock the encrypted file, but as a couple of the earlier mentioned product names suggest, you only need to remember a single one.

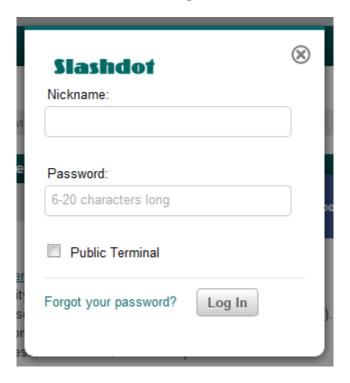
Here's the critical point: *this single password must be strong*! If you're going to lock up the keys to every single website with just one password, you can forget about birthdays and kids names and sandwiches, you really need to pick

something decent this time.

The 1Password approach

Running 1Password, let me show you what happens when I log on to a website in the traditional way. I'm going to log into <u>Slashdot</u> which is a bit of a techie website but the process is pretty much the same for almost every website out there.

We start off with the usual username and password:



But after I hit the "Log In" button, 1Password offers to save the credentials:

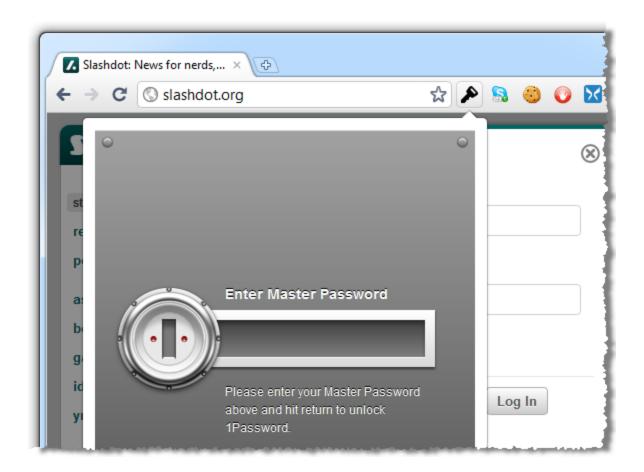


The name defaults to the address of the page but I can always rename it to something more logical either now or a little later on. Once I hit the "Save" button, 1Password asks me for the "Master" Password", that is the single password required to manage all my other ones:

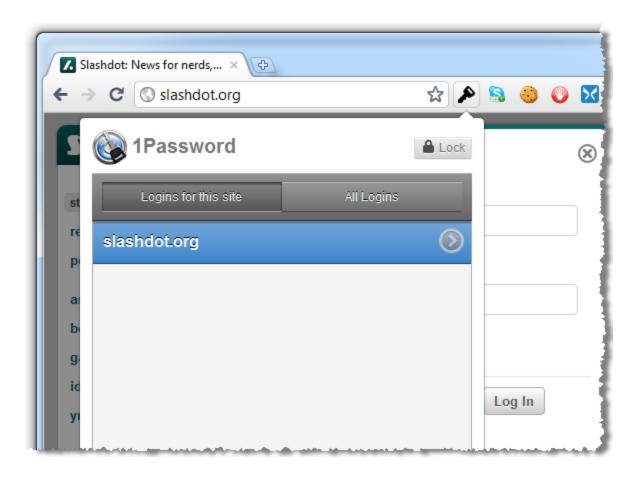


This is one, single, strong password which I have memorised. In fact it's now the *only* one I've memorised and no, it's not "Iloves@nDwich3s"!

With this saved, let me now log out of Slashdot then go back and attempt to login again but this time, rather than entering my Slashdot credentials (which I've conveniently and deliberately forgotten), I'm going to hit the little key icon to the right of the URL bar:



This is now asking for my master password again – the only one I ever need to remember. After entering this, I can see the entry created earlier on:



I *could* have multiple entries in here (you might have more than one account at a particular site), but I'll just double click on the existing entry. And that's it – we're now logged on!

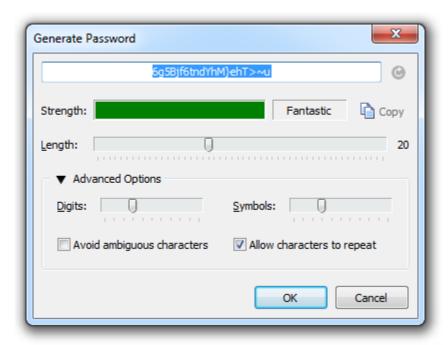
The beauty of this process is that it's identical for every single site. I don't need to remember those 90 odd passwords any more, I simply need to go through the motions of manually logging onto each site once and allowing 1Password to save the credentials.

You can also do this from different browsers. I'm using Google Chrome in the examples above but 1Password also integrates with other browsers.

Getting secure

Of course the chances are your passwords aren't real secure to begin with and all this process is doing is keeping a secure record of bad passwords. This is a great time to do some housekeeping and 1Password makes it very easy.

When I went through and added all my accounts, each time I came across one with a weak password I went into the 1Password application, opened up the account I just created and generated a new one. There's a really neat little tool built right in which makes this a breeze:



This is what a secure password looks like (highlighted in blue above). If it's not something you need to be a savant to memorise, it's not secure enough. But of course with the process described above it doesn't matter that the password is entirely unintelligible, all you need to remember is your master password.

Now, this process won't actually change your password on the website, only the one you have recorded in 1Password. You'll need to copy this one into your clipboard then go onto the individual website and change it accordingly. Yes, it's a bit of mucking around but for the sake of a few minutes you've just created a very secure, very unique password which can't be used against you on any of your other online accounts.

There's one gotcha in all of this; some websites don't let you create secure passwords. Earlier this year I wrote about the Who's who of bad password practices – banks, airlines and more where I found that some websites – especially banks, oddly enough – simply won't let you construct long, random passwords. Either they limit the length to a very low number, they disallow many character types or in extreme examples, they insist on a short PIN containing only numbers. Unfortunately you're entirely at the mercy of the controls these sites place on passwords so when you hit a limitation like this all you can do is maximise what you can within a ridiculous constraint.

Taking your passwords with you

One thing that was important to me was that I could access my passwords from any location, on any device, at any time. Work PC, home PC, iPad and iPhone all needed to sync up.

1Password lets you do all of this by using the <u>Dropbox</u> file syncing service. This is a great product which has proven very robust and is easy to configure to keep your 1Password file synced. In the end, it means all my PCs have the same secure password file and my iPad and iPhone respectively have friendly little apps like these:



Is it risky putting the password file online? Well there's a *degree* of risk, sure, but the Dropbox service has proven a very secure implementation over the years. And of course the 1Password file is still securely encrypted so even if someone gets their hands on it, they still need the (strong) master password. In fact the weakest link in the whole thing is probably the password you secure your Dropbox account with which, by now of course, is also very strong:)

Isn't this "all your eggs in one basket" stuff?

Yes, it is, but it's a basket that is very well thought out and very firmly secured. Someone would have to firstly obtain the file containing all the passwords exposed and secondly have your master password either disclosed, guessed or brute force attacked, none of which should happen if you choose one securely.

Whilst having all your account details exposed at once is undoubtedly a very bad

thing, the risk is infinitesimal compared to the chances of having it breached via website.

Of course the other risk is that an as yet unknown vulnerability is found with the 1Password software. Certainly what we'd call a <u>zero-day</u> vulnerability (one that is not yet known), is possible. In fact there was <u>one found in LastPass just last month</u> and to their credit, they plugged that hole in no more than a few hours. And that's the point with professional products of this nature; their entire being is centred on offering a secure solution and if a vulnerability is found, you can be pretty damn sure it's going to be squashed *very* quickly.

Summary

So now that you've got all this super security, you're pretty much invincible right? Uh...

And this brings me to a neat philosophical conclusion; security is all about risk mitigation -you never actually become "secure", you merely decrease your risk. On balance, the risk of your account details sitting out there in even a very secure website is significantly higher than having them sit there in your 1Password file.

But beyond just security, the password manager route is a very handy solution. Having all your accounts handy on all your devices and being able to simply logon with the once strong password is a very convenient route indeed.

And finally, when the time comes that you realise one of your accounts has been breached (and trust me, *it will come*), it's no good thinking about password security then – it's too late. So put aside a few hours one afternoon, spend just a few dollars and get yourself organised. Either that or start developing a taste for acai berries!

Comments

It's a pet peeve of mine when the weak passwords in the Gawker database are used as some kind of evidence that people are bad at making strong passwords. I had an account there, with my throwaway password, which wasn't on your top 25 list but is a single English word.

Why did I choose such a bad password? Because I don't care about my Gawker account. So why should I choose a password I can't remember, encrypt it in a strong file on my dropbox and then have to go get it every time I want to write a semi-anonymous comment?

I would guess that at least 75% percent of the top 25 list's uses are from throwaway accounts made because Gawker makes everyone who comments have them.

Heck, Bruce Schneier recommends using one easy throwaway password for sites you don't care about. Seems like good advice to me. Good luck brute forcing my bank password, but Gawker... you can have it for free if you like.

Fantastic! Now I can completely forget about trying to remember passwords and simply use a 1 password style application when I try to visit a website from my wp7 phone, internet-connected tv, video game console, friend's computer... This definitively and completely solves the problem once and for all.

The author demonstrates a laughable lack of understanding of what is secure in a password. You don't need to add numbers, two cases, and punctuation for a password to be secure. You don't need it to not be easily remembered words.

In fact, a password made of four common, easily-remembered English words is more secure than an eight character password with upper, lower, number, and punctuation.

So all you need to do is come up with four random words that you can visualize in an amusing way, the common memorization technique, and you have an easy to remember, secure password.

No uppercase

No numbers

No punctuation

All easy, dictionary words

Troy: The author's suggestion is that your advice is perfectly legitimate... when you only have a few passwords. And they allow spaces. And you're not limited to a low character count. And it doesn't demand a PIN. For the real world where we are faced with more accounts than we can remember, often onerous password restrictions and frequently accounts we may only authenticate to a few times a year - beyond the memory span of individual collections of words - there are password managers.

What the link you provide boils down to, along with the original post above, is "we're finding ways to justify that you spend money on this completely unnecessary software that does something you can do with freeware, or in other ways, without it".

You present needlessly elaborate examples of what we're supposedly facing, ignoring (for example) how mnemonic techniques like forming a funny mental image of four words works across an effectively unlimited list of items. A salesman who is terrible with names can end up remembering the names of hundreds, even thousands, of people with similar, simple techniques...and salesmen don't tend to be the brightest human beings.

In the unlikely event that someone needs a password manager, they can simply google "password manager", and find many others for less or free, without the deceptive hype overstating both risks and difficulty like the above nonsense.

Troy: I went to this restaurant the other night and the waitress goes around the table taking orders from half a dozen people. Entrees, mains, how they want their steak done, what they wanted to drink etc. All committed to memory, all

delivered verbatim. We all marvelled at her ability to do this but upon reflection, perhaps you're just one of those people with a waitress-like memory. For the rest of us, there are many good password managers, free or commercial, but that seems to be your conclusion anyway. They have a place, perhaps just not for you.

Epilogue

Looking back at this post more than a decade on, there are so many good things it's done for me. First of all, it really helped me carve out a niche as someone who gives a lot of thought to how authentication schemes work. I'm loath to say it helped make me an "expert" because I always find that an overly self-ingratiating term, but it certainly helped me make a name for myself. It also gave me a really good opportunity to start thinking about the human side of security, a theme that would prevail across so many subsequent blog posts. The fact that people defer to the simplest possible path around technology barriers such as password complexity criteria meant that relying on humans alone to create decent passwords was simply never going to work.

The 1Password relationship also became extremely valuable over the years. The folks there saw the post and we began conversing, often just on the public timeline but it also gave me the opportunity to start building out closer relationships with the people there. There was nothing of a commercial nature in those early days and that was for the best; nobody ever questioned my motives or suggested that I was somehow incentivised to align myself to the password manager I so frequently promoted. Eventually though, that changed as I entered into an agreement with them to have product placement on Have I Been Pwned a full 8 years after the original blog post. Another 2 and a half years after that, I joined their board of advisors.

There's a very profound observation I want to make here that's enormously important to me personally: I've done very well financially out of this industry, but every single thing I can think of that's in my commercial best interest today, began as a community-first effort without there ever being an expectation of monetary reward. I had no idea I'd team up with 1Password in any way whatsoever when I wrote this blog post, I just wanted to share my thoughts on how people should manage their passwords. I couldn't have cared less about turning that content into dollars at the time and I believe - I'm convinced - that this is one of the primary contributors to my later success. It was never about the money, it was about the content, but you put enough good stuff out there that's valuable enough to people and financial opportunities inevitably follow.

A BRIEF SONY PASSWORD ANALYSIS

As 2011 marched on, I was finding myself drawn into the world of data breaches and passwords. What I found particularly fascinating was that in data breaches, you have someone else's dirty laundry on display for everyone to see. Not intentionally either - not like with open-source software - rather this was a "warts and" all look at how online systems had been put together and oh boy, there were some shockers!

But breaches didn't just lay the design decisions of the company involved out for everyone to see, they also exposed the security decisions of their customers by virtue of the passwords they'd chosen. What I was finding fascinating here is even though we knew anecdotally that people generally create short, weak passwords, when data breaches expose them in plain text, we can establish that *empirically*. I'm a numbers guy and it meant a lot to me to be able to take the passwords from Sony and plot them into graphs, demonstrating precisely how bad humans really are at creating passwords.

06 JUNE 2011

analysed Sony Pictures passwords after their breach, a practice that would later lead me to create HIBP...

So the Sony saga continues. As if the whole thing about <u>77 million breached</u> <u>PlayStation Network accounts</u> wasn't bad enough, <u>numerous other security</u> <u>breaches</u> in other Sony services have followed in the ensuing weeks, most recently with <u>SonyPictures.com</u>.

As bad guys often like to do, the culprits quickly stood up and put their

handiwork on show. This time around it was a group going by the name of LulzSec. Here's the interesting bit:

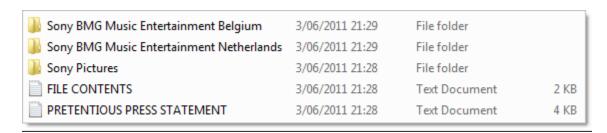
Sony stored over 1,000,000 passwords of its customers in plaintext

Well actually, the really interesting bit is that they created a <u>torrent of some of the breached accounts</u> so that anyone could go and grab a copy. Ouch. Remember these are innocent customers' usernames and passwords so we're talking pretty serious data here. There's no need to delve into everything Sony did wrong here, that's both mostly obvious and not the objective of this post.

I thought it would be interesting to take a look at password practices from a real data source. I spend a bit of time <u>writing about how people and software manage passwords</u> and often talk about things like entropy and reuse, but are these really discussion worthy topics? I mean do people generally get passwords right anyway and regularly use long, random, unique strings? We've got the data – let's find out.

What's in the torrent

The Sony Pictures torrent contains a number of text files with breached information and a few instructions:



The interesting bits are in the "Sony Pictures" folder and in particular, three files with a whole bunch of accounts in them:

Sony_Pictures_International_AUTOTRADER_USERS	3/06/2011 21:29	Text Document	1,645 KB
Sony_Pictures_International_BEAUTY_USERS	3/06/2011 21:29	Text Document	674 KB
Sony_Pictures_International_COUPONS	3/06/2011 21:29	Text Document	452 KB
Sony_Pictures_International_DELBOCA_USERS	3/06/2011 21:29	Text Document	573 KB
Sony_Pictures_International_MUSIC_CODES	3/06/2011 21:29	Text Document	1,369 KB
Sony_Pictures_International_TABLE_LAYOUT	3/06/2011 21:28	Text Document	28 KB

After a little bit of cleansing, de-duping and an import into SQL Server for analysis, we end up with a total of 37,608 accounts. The LulzSec post earlier on did mention this was only a subset of the million they managed to obtain but it should be sufficient for our purposes here today.

Analysis

Here's what I'm really interested in:

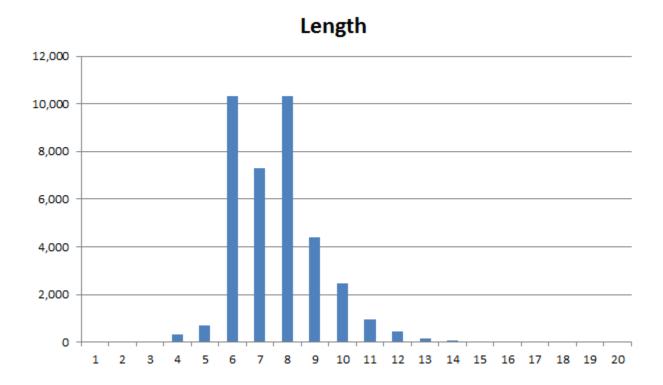
- 1.Length
- 2. Variety of character types
- 3.Randomness
- 4. Uniqueness

These are pretty well accepted measures for password entropy and the more you have of each, the better. Preferably heaps of all of them.

Length

Firstly there's length; the accepted principle is that as length increases, as does entropy. Longer password = stronger password (all things else being equal). How long is long enough? Well, part of the problem is that there's no consensus

and you end up with all sorts of opinions on the subject. Considering the usability versus security balance, around eight characters plus is a pretty generally accepted yardstick. Let's see the Sony breakdown:



We end up with 93% of accounts being between 6 and 10 characters long which is pretty predictable. Bang on 50% of these are less than eight characters. It's interesting that seven character long passwords are a bit of an outlier – odd number discrimination, perhaps?

I ended up grouping the instances of 20 or more characters together – there are literally only a small handful of them. In fact there's really only a handful from the teens onwards so what we'd consider is a relatively secure length really just doesn't feature.

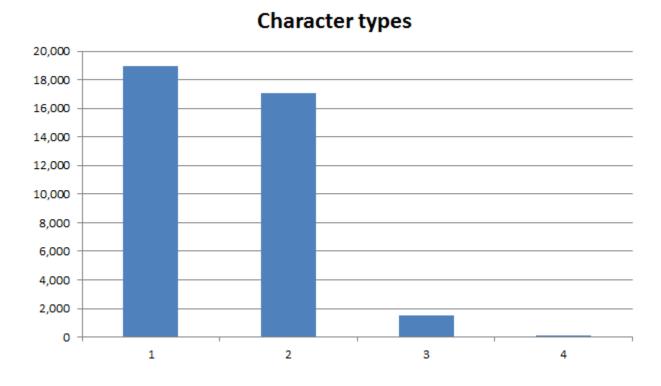
Character types

Length only gives us so much, what's really important is the diversity within

that length. Let's take a look at character types and we'll categorise them as follows:

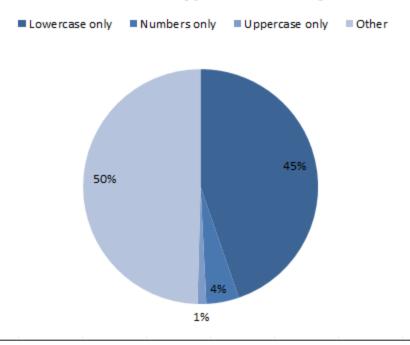
- 1.Numbers
- 2.Uppercase
- 3.Lowercase
- 4. Everything else

Again, we've got this issue of usability and security to consider but good practice would normally be considered as having three or more character types. Let's see what we've got:



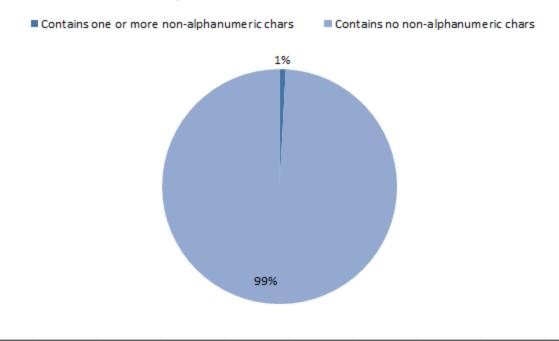
Or put another way, only 4% of passwords had three or more character types. But it's the spread of character types which is also interesting, particularly when only a single type is used:

Character type exclusivity



In short, half of the passwords had only one character type and nine out of ten of those where all lowercase. But the really startling bit is the use of nonalphanumeric or characters:

Alphanumeric characters



Yep, less than 1% of passwords contained a non-alphanumeric character. Interestingly, this also reconciles with the <u>analysis done on the Gawker database</u> a little while back.

Randomness

So how about randomness? Well, one way to look at this is how many of the passwords are identical. The top 25 were:

seinfeld, password, winner, 123456, purple, sweeps, contest, princess, maggie, 9452, peanut, shadow, ginger, michael, buster, sunshine, tigger, cookie, george, summer, taylor, bosco, abc123, ashley, bailey

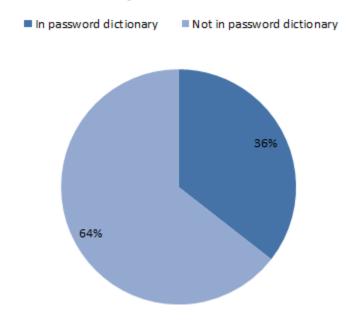
Many of the usual culprits are in there; "password", "123456" and "abc123". We saw all these back in the <u>top 25 from the Gawker breach</u>. We also see lots of passwords related to the fact this database was apparently related to a competition: "winner", "sweeps" and "contest". A few of these look very specific

(9452, for example), but there may have been context to this in the signup process which lead multiple people to choose the same password.

However in the grand scheme of things, there weren't a whole lot of instances of multiple people choosing the same password, in fact the 25 above boiled down to only 2.5%. Furthermore, 80% of passwords actually only occurred once so whilst poor password entropy is looking rampant, most people are making these poor choices independently and achieving different results.

Another way of assessing the randomness is to compare the passwords to a password dictionary. Now this doesn't necessarily mean an English dictionary in the way we know it, rather it's a collection of words which *may* be used as passwords so you'll get things like obfuscated characters and letter / number combinations. I'll use this one which has about 1.7 million entries. Let's see how many of the Sony passwords are in there:

Prevalence of password in dictionaries



So more than one third of passwords conform to a relatively predictable pattern. That's not to say they're not long enough or don't contain sufficient character types, in fact the passwords "1qazZAQ!" and "dallascowboys" were both matched

so you've got four character types (even with a special character) and then a 13 character long password respectively. The thing is that they're simply not random – they've obviously made appearances in password databases before.

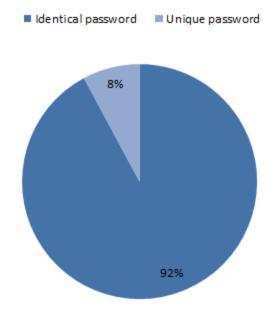
Uniqueness

This is the one that gets really interesting as it asks the question "are people creating unique passwords across multiple accounts?" The thing about this latest Sony exploit is that it included data from multiple apparently independent locations within the organisation and as we saw earlier on, the dump LulzSec provided consists of several different data sources.

Of particular interest in those data sources are the "Beauty" and "Delboca" files as they contain almost all the accounts with a pretty even split between them. They also contain well over 2,000 accounts with the same email address, i.e. someone has registered on both databases.

So how rampant is password reuse between these two systems? Let's take a look:

Password reuse

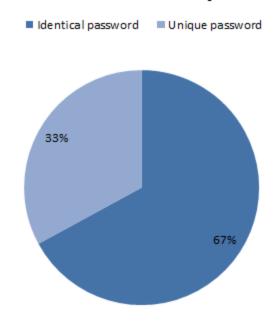


92% of passwords were reused across both systems. That's a pretty damning indictment of the whole "unique password" mantra. Is the situation really this bad? Or are the figures skewed by folks perhaps thinking "Sony is Sony" and being a little relaxed with their reuse?

Let's make it really interesting and compare accounts against Gawker. The internet being what it is there will always be the full Gawker database floating around out there and a quick Google search <u>easily discovers live torrents</u>. Gnosis (the group behind the Gawker breach) was a bit more generous than LulzSec and provided over 188,000 accounts for us to take a look at.

Although there were only 88 email addresses found in common with Sony (I had thought it might be a bit higher but then again, they're pretty independent fields), the results are still very interesting:

Password reuse across Sony and Gawker



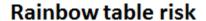
Two thirds of people with accounts at both Sony and Gawker reused their passwords. Now I'm not sure how much crossover there was timeframe wise in terms of when the Gawker accounts were created versus when the Sony ones were. It's quite possible the Sony accounts came after the Gawker breach (remember this was six months ago now), and people got a little wise to the non-unique risk. But whichever way you look at it, there's an awful lot of reuse going on here.

What really strikes me in this case is that between these two systems we have a couple of hundred thousand email addresses, usernames (the Gawker dump included these) and passwords. Based on the finding above, there's a statistically good chance that the majority of them will work with other websites. How many Gmail or eBay or Facebook accounts are we holding the keys to here? And of course "we" is a bit misleading because anyone can grab these off the net right now. Scary stuff.

Putting it in an exploit context

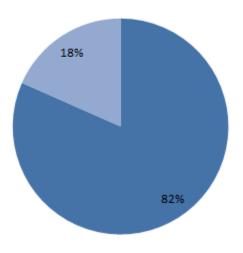
When an entire database is compromised and all the passwords are just sitting there in plain text, the only thing saving customers of the service is their password uniqueness. Forget about rainbow tables and brute force – we'll come back to that – the one thing which stops the problem becoming any worse for them is that it's the only place those credentials appear. Of course we know that both from the findings above and many other online examples, password reuse is the norm rather than the exception.

But what if the passwords in the database were hashed? Not even salted, just hashed? How vulnerable would the passwords have been to a garden variety rainbow attack? It's pretty easy to get your hands on a rainbow table of hashed passwords containing between one and nine lowercase and numeric characters (RainbowCrack is a good place to start), so how many of the Sony passwords would easily fall?









82% of passwords would easily fall to a basic rainbow table attack. Not good, but you can see why the rainbow table approach can be so effective, not so much because of its ability to make smart use of the time-memory trade-off scenario, but

simply because it only needs to work against a narrow character set of very limited length to achieve a high success rate.

And if the passwords were salted before the hash is applied? Well, more than a third of the passwords were easily found in a common dictionary so it's just a matter of having the compute power to brute force them and repeat the salt plus hash process. It may not be a trivial exercise, but there's a very high probability of a significant portion of the passwords being exposed.

Summary

None of this is overly surprising, although it remains alarming. We know passwords are too short, too simple, too predictable and too much like the other ones the individual has created in other locations. The bit which did take me back a bit was the extent to which passwords conformed to very predictable patterns, namely only using alphanumeric character, being 10 characters or less and having a much better than average chance of being the same as other passwords the user has created on totally independent systems.

Sony has clearly screwed up big time here, no doubt. The usual process with these exploits is to berate the responsible organisation for only using MD5 or because they didn't salt the password before hashing, but to not even *attempt* to obfuscate passwords and simply store them in the clear? Wow.

But the bigger story here, at least to my eye, is that users continue to apply lousy password practices. Sony's breach is Sony's fault, no doubt, but a whole bunch of people have made the situation far worse than it needs to be through reuse. Next week when another Sony database is exposed (it's a pretty safe bet based on recent form), even if an attempt has been made to secure passwords, there's a damn good chance a significant portion of them will be exposed anyway. And that is simply the fault of the end users.

Conclusion? Well, I'll simply draw back to a previous post and say it again: <u>The only secure password is the one you can't remember.</u>

Update, 7 June 2011:

I've cleaned up a couple of the graphs and would also like to clarify some of the points coming through in the comments:

- 1. There does not appear to have been either length or character type restrictions on the databases. There are lengths ranging from 1 to 35 characters and (occasional) uses of non-alphanumerics.
- 2.As far as I know, this database is not directly related to PSN. Sony is a huge media network and LulzSec claims this came from SonyPictures.com. There is no evidence to suggest that these passwords were created using a handheld game console controller.

Comments

Until engineers learn to understand human nature, this will continue to happen.

I cannot believe commenters (and the writer) are actually still recommending that people adopt even harder to remember paswords. Engineers need to get their heads out of their calculators and see people for what they are.

With the vast number of websites people use, it is IMPOSSIBLE to have a unique password for each one. Password re-use is simply a natural consequence of using the web. Secondly, complex passwords are hard to remember, and people want ease of use and convenience.

Who wants to (or is even able to) remember twenty passwords that look like

"sdfgh*7&456#56?7DGBFD"?

Thirdly, people are pretty good at managing risk vs stress trade-offs. The probability of being hacked is low, and the potential loss is low, so why should they stress themselves remembering multiple difficult passwords? Not worth it.

If engineers and security specialists are really so excited about improving security, they should realise they can't change human nature. So instead of berating people for using weak passwords, they should realise the fault is with them for dsigning a system that is not human friendly in the first place.

Troy: It's impossible to remember strong, unique passwords for every site but it's not impossible to implement them. In fact it's very simple - use a password manager: http://www.troyhunt.com/201...

you CAN however create 1 or 2 strong passwords that have say, the first 3 characters "reserved" for a site specific string.

for example amz21\$Flexi%! = amazon cit21\$Flexi%! = citibank

etc. This way the pword is both specific to the site and rememberable (assuming that 21\$Flexi%! means something to you.)

It would not be easily hacked unless someone cross referenced hacked databases. If you split it up to 3 strong passwords with flexible strings, you decrease your chances of corrolation.

And then when your own computer is compromised they get *all* of your passwords and will know every service which to use them to access. Nice.

Troy: It doesn't quite work that way Paul, good password managers like

1Password store all passwords in a cryptographically secure fashion. Accessing the keychain is worthless without the master password which should not be accessible even if the machine is owned.

Right, but that does nothing to keep someone from physically (or legally) coercing me into revealing a master password to unlock my password database and all corresponding systems I have access to. Granted, the alternative is impossible, to remember secure-enough passwords to all the systems I access, but at least this should be acknowledged as a potential concern.

Troy: Sure, but that's a far cry from "when your own computer is compromised they get *all* of your passwords".

Look, you're not going to get a 100% "secure" solution so it's a matter of balance. I would argue that the likelihood of someone first owning my machine then grabbing my 1Password keychain and then coercing me into disclosing my master password is infinitesimally small compared to a website where I use a set of credentials being compromised.

Epilogue

I think what's most fascinating about looking back on this post is that a decade on, nothing has really changed. People still make terrible password choices, and they still reuse them all over the place. 2FA? Barely used. Passwordless solutions? We've got more passwords now than ever. Plus, we've got twice as many people on the internet now than we did when I first wrote this post so all that's happened is, we've got more people making the same bad decisions!

This blog post would later prove to be a big part of the foundation for building Have I Been Pwned 2 and a half years later. In fact, in the launch blog post I link to the Sony password piece in the second sentence, citing it as one of the inspirations for building HIBP. Just like I commented with the 1Password epilogue earlier, this is a really key observation for how my career has unfolded; I didn't write the Sony passwords post with any expectations whatsoever other than to scratch an itch by sharing something I thought was interesting. That an opportunity later came as a result of that blog post was purely accidental, a pattern that would become a recurring theme in my career.

FIND MY CAR, FIND YOUR CAR, FIND EVERYBODY'S CAR; THE WESTFIELD'S IPHONE APP PRIVACY SMORGASBORD

Ah, my first public smackdown of a big corp! This blog post was followed by many more that shone a light on egregiously bad security problems, but it all started here with this one. It began the same way as many of the others that would follow also did: with curiosity. Time and time again, I'd look at a system and think to myself: "I wonder what would happen if I just did [thing]". It's the same for so many other folks in this industry as well where that curiosity drives us to poke and prod and see if we can get a system to do something it was never designed to. None of this was done with malicious intent, of course, it just followed the hacker ethos of wanting to explore the inner workings of a service. And then break it

14 SEPTEMBER 2011

he Westfield parking privacy debacle was the first time I really went to town on a company's security, and also the first time I was actually worried about the potential repercussions.

When news came through recently about the Bondi Westfield shopping centre's new "Find my car" feature, the security and privacy implications almost jumped off the page:

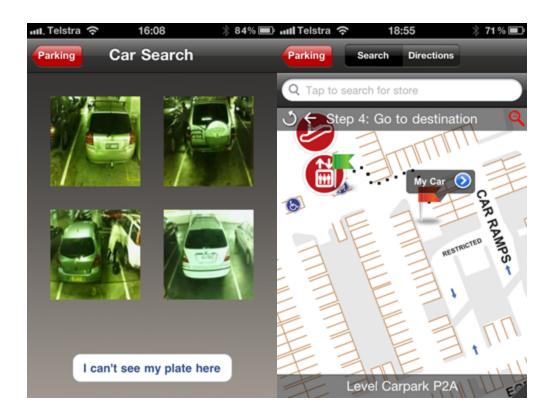
"Wait – so you mean all I do is enter a number plate – any number plate – and I get back all this info about other cars parked in the centre? Whoa."

If that statement sounds a bit liberal, read on and you'll see just how much information Westfield is intentionally disclosing to the public.

Intended use

Let's begin with how the app looks to the end user. This all starts out life as the Westfield malls app in the iTunes app store and for some time now, it's been able to help you find stores in the centre. As of recently though, it has a "Parking" feature which allows you to enter a number plate and get back a series of images then receive directions on how to navigate to the one which appears to be your vehicle. Perhaps Westfield drew inspiration from Seinfeld's The Parking Garage on this one! Here's how it all ties together:





To the casual user of the application, the number plates – and this is what I'm really talking about when I say "privacy" – appear to be indiscernible. Certainly it's not clear from the images above but it's also not clear after screen grabbing and expanding it:



The number plate is actually AWC11A, but we'll get back to that.

Anyway, this is all made possible by using the <u>Park Assist</u> technology which puts the little guy in the image below on the roof between each park so they can both notify customers of vacant spots and snap pictures of them once they park:



The interesting bit though is that the implementation of this app readily exposes some fairly serious, rather extensive data that many people would probably be concerned about. *And* it doesn't have to.

Under the covers

The way these smart phone apps tend to work is that when they have a dependency on external data retrieved from the internet, is they communicate backwards and forwards via services which travel over the same protocol as most of your other internet traffic – HTTP. Very often these services contain all sorts of information with only a small subset actually being exposed to the user via the application consuming the service. In Westfield's case it was fair to assume that this service would contain some information about the vehicles matching the number plate search and what their location is.

Using a free tool like <u>Fiddler</u> and allowing it to <u>act as an HTTP gateway for the iPhone</u>, it's easy to interpret and inspect the contents of communication between the app and the server it's talking to. When I did this for the Westfield app, here's what I found:

#	Result	Protocol	Host	URL	Body	Caching	Content-Type
{ ^{js} } 1	200	HTTP	120.151.59.193	/v2/bays.json?	1,476	Expires: Tue	application/json; charset=utf-
2	200	HTTP	120.151.59.193	/v2/camera-se	4,068		image/jpeg
🛐 3	200	HTTP	120.151.59.193	/v2/camera-se	4,078		image/jpeg
9 4	200	HTTP	120.151.59.193	/v2/camera-se	4,310		image/jpeg
🛐 5	200	HTTP	120.151.59.193	/v2/camera-se	3,792		image/jpeg

What we're seeing here is a total of five requests made to Westfield's server: The first one returns a <u>JSON</u> response which contains the data explaining the location of cars matching the search. The next four requests are for images which are pictures of the cars returned by the search. Here's what we get:



Apart from the slight difference in aspect ratio, this is exactly what we saw in the original app so no surprises yet. But here's where it gets really interesting – let's examine that JSON response. Firstly, it's a GET request to the following address:

<u>h t t p : / / 1 2 0 . 1 5 1 . 5 9 . 1 9 3 / v 2 / b a y s . j s o n ?</u> <u>visit.plate.text=abc123~0.3&is_occupied=true&limit=4&order=-similarity</u>

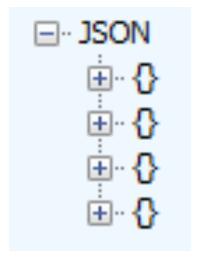
One of the nice things about a <u>RESTful</u> service like this is the ability to easily pass parameters in the request. In the URL above, we can see four parameters:

- 1.The number plate we're searching for appended with " \sim 0.3"
- 2.An "is_occupied" value set to "true"

3.A "limit" set to "4"

4.An "order" set to "-similarity"

Now when we look at the response body, we see the following:



What this is telling us is that the JSON response contains four collections of data. Let's expand that first one and see what's inside:

```
— group

    ...id=221
    --- name = P3A. 13
   :...sign_offset=0
  ·id=37210
 - is_in_violation=False
  -is_occupied=True
  - is_out_of_service = False
...map
    ...id=5
   ....name=P3A.png
... policy
    ...id=1
   .... name=No Limit
...x=769
   .... v=546
--- address=10.8.14.137-2-10
    ··· confidence = 100

<u>□</u> · extended_properties

       --- mode =0
       --- polygon_timestamp=2011-09-13T15:53:17.9661394+10:00
       ...position=10
        ...row=2
       --- row_controller = 10.8.14.137
       --- software_version=v02.016.465 Sep 1 2010
       — thumbnail_request_timestamp=2011-09-13T15:33:11.8638795+10:00
        -- thumbnail_timestamp=2011-09-13T15:33:20.5205727+10:00
       ··· thumbnail_update_reason=TooOld
       thumbnail_update_reason_timestamp=2011-09-13T05:33:11.0000000+00:00
     -last_contact=2011-09-13T15:53:11.5127253+10:00
     -- plugin=Serial
Ė∴ type
    ...id=1
     name=Casual
Ė∵visit
     dwell=01:13:50.7892294
     entry_timestamp=2011-09-13T04:43:15.0000000+00:00
     -id=28286945
   ... plate
        -confidence=97
         text=AWC11A
         timestamp=2011-09-13T14:45:23.6061075+10:00
_ zone
    ...id=3
     name=P3
```

This is a fair bit of data. Actually it's *a lot* of data and it's being sent down to your phone every time you try to locate a car. Remember, all the app needs to do is show us an image of what may be our car. But the really worrying bit is what's

inside the "visit" node; Westfield is storing and making publicly accessible the time of entry and the number plate (see the "text" field) of what appears to be every single vehicle in the centre. What's more, it's available as a nice little service easily consumable by anyone with the knowhow to build some basic software.

But this is only four results, right? Actually, it's worse than that. A lot worse. That URL for the service endpoint we looked at earlier contains a number of parameters – filters, if you like – and removing these readily provides the current status of all 2,550 sensors. This includes the number plate of any car currently occupying a space and as you can see, *it's available by design to anyone*:

http://120.151.59.193/v2/bays.json

You can freely request that resource over and over as many times as desired and then store the data to your heart's content. Now *that*, is a privacy concern.

The impact to privacy

What this means is that anyone with some rudimentary programming knowledge can track the comings and goings of every single vehicle in one of the country's busiest shopping centres. In an age where we've become surrounded by surveillance cameras we expect our movements to be monitored by the likes of centre management or security forces, but not on public display to anyone with an internet connection!

Think about the potential malicious uses if you're able to write a simple bit of software:

- 1.A stalker receives a notification when their victim enters the car park (and they'll know exactly where the victim is parked).
- 2.A suspicious husband tracks when his wife arrives and then leaves the car park.

- 3.An aggrieved driver holding a grudge from a nearby road rage incident monitors for the arrival of the other party.
- 4.A car thief with their eye on a particular vehicle could be notified once it is left unattended in the car park.

With Westfield standing up the service in the way they have, *this becomes* extremely easy. Furthermore, this is just one shopping centre out of dozens of Westfields across the country. If this practice continues, data mining the movements of individual vehicles across shopping centres will be a breeze *for* anyone with basic programming knowledge. And that's really the crux of the problem in that this isn't one of those "Oh no, the big corporation is tracking me" situations, it's that anyone can track me.

Whilst I'm by no means a strong privacy advocate (I have a fairly open life on display through numerous channels on the web), something about this just doesn't sit quite right with me. Certainly those people who *are* strong privacy advocates would object to such a public disclosure of information.

What needs to be done

Putting my "software architect / security hat" back on for a moment, the problem is simply that Westfield is exposing data this application has no need for. The best way to keep a secret is to never have it and this is where they've gone wrong.

The parking feature of the app is designed for only one purpose: taking a number plate from the user and returning four possible positions with grainy images of the vehicle. On this basis, every piece of data in the "visit" node in the image earlier on is totally unnecessary, as is the ability to pull back more than four records at a time and as is the ability to do it over and over again as fast as

possible. All that is required is the image so that someone can visually verify it's their car (the number plate need not be clear), and of course information on the location within the centre.

If they were to do this, the privacy risk is dramatically reduced as all you're left with now as Joe Public is a small bunch of grainy images with indecipherable number plates. The positive feedback of the service explicitly returning the number plate (and degree of confidence in its integrity), is gone. Sure, there's still a privacy risk in that I can manually open up the app and search for someone's car then manually ID it, but the potential for automation is gone.

In fact most of the data returned in that service is totally unnecessary. Trimming it back would not only (largely) resolve the privacy problem, it would also reduce the size of the service hence speeding it up for the end user and reducing the bandwidth burden on Westfield. Win-win-win.

Summary

In the process of researching and writing this post I also identified other major vulnerabilities of a rather serious nature. I've done the right thing and attempted to notify Westfield of these and won't publicly disclose them.

The information above, on the other hand, is already public knowledge in so far as people know there is a database containing their cars that is publicly accessible, they just probably don't know how easy it is to get hold of. But of course *by design*, this information is intended to be consumed on demand, as frequently as possible and for any vehicle in the car park.

All in all, this just doesn't seem to be very well thought out on behalf of the developers. In fact on that basis, if Park Assist is behind this and they've implemented the same system in other locations, the Bondi situation could just be the tip of the iceberg. In all likelihood, it's not Westfield's intent to expose

this volume of information in such an easily consumable fashion and hopefully they'll ask for the software to be revised once they're aware of the full extent of the situation.

Update 1, 11:10 Sep 14:

A couple of hours after posting this, a helpful reader contacted Park Assist directly and they responded promptly by pulling the service altogether (they appear to control the environment). Based on their response, it seems the API into the service should never have been publicly exposed and used by the iPhone app in this fashion. It also appears that in addition to the privacy risk, further security vulnerabilities were identified by other individuals. Whilst I won't disclose these publicly, I did attempt to contact Westfield directly about the issue. As of now, I'm still to get a response from them.

Update 2, 12:00 Sep 14:

I just had a good phone chat with the guy from Park Assist who provided the response linked to in the previous update. In short, they've handled this situation very efficiently and have responded with urgency and professionalism. We discussed some general security concepts including how the service could be properly secured for future use and it's clear they understand precisely what needs to be done. It's unfortunate for them that the software was configured in this fashion in the first place but certainly they're doing everything right in their response.

Comments

My name is John Batistich and I am the General Manager Marketing for Westfield. Firstly, thanks Troy for bringing this issue to our attention. Our intention was to create a free service for our customers so they never lose their car again! However, we have more work to Our partner, Park Assist, who provide the camera technology to capture the number plate today advised there was an issue with the authentication of their data feed to the iPhone which resulted in number plate data being publicly assessable via the internet. This issue has been addressed immediately by Park Assist and the Find My Car functionality will not be available for approximately one week until the app has been modified to ensure that data cannot be publicly assessable online. Further, the 'Find my Car' functionality on our app is similar to other location-based services and has been developed to provide a service to the average shopper, in an effort to make it easier to find their car. In terms of privacy, the application does not contravene the Privacy Act in so far as numbers plates are not "personal information", and are therefore not subject to that act. Having said that, the application theoretically could be used for purposes other than its original intention, however it does not facilitate any activity that couldn't already happen otherwise. For example, a member of the general public may try to use the application to find a car that is not theirs. On the other hand, at the request of police, the application might also be used to assist in their enquiries into a given situation however, Westfield would not expect either of these situations to be typical. We appreciate you bringing this issue to our attention and we are now working on an update to resolve the technical issue.

Troy: Thanks for posting John, I'm sure this came as a big surprise to you. Certainly the service can still be provided to the end user without any change in the overall experience to them, it's simply the underlying implementation which poses the privacy risk.

The activity which could be performed under the previous implementation which could not happen otherwise is the positive and automated identification of every vehicle in the car park. The fact that any anonymous member of the public could get a precise list of every car in a matter of seconds without actually knowing any valid number plate - and then continuously repeat the process -

simply wasn't possible without your application. Conversely, once the software is rectified, the extent of discovery will simply be grainy images of vehicles with illegible number plates which closely match the pattern entered into the application. There's a world of difference in these two scenarios.

As I said in Update 2, Park Assist appear to have acted promptly and certainly from what I've observed, have taken the appropriate action. I've offered support in testing if desired and I hope you can rectify and relaunch the service promptly.

Nice work.

Did you contact Westfield or parkassist?

Wouldn't some kind of responsible disclosure have been better than putting this exploit out in the wild?

Troy: Hi Raphael, yes, as per the summary I have tried to contact Westfield about other security related vulnerabilities that I won't publicly discuss. The info above is purely about the privacy impact of the service intended for public consumption.

Oh, sorry, too much text :)

Nevertheless, what were the reasons for writing about the privacy vulnerabilty while it is still exploitable?

Troy: The reason is simply for awareness. Consumers reasonably expect to know the extent to which they're being monitored and how public that information is

made and that isn't being made apparent to them. The point I'm making above is that the ability to monitor this is by design; this is how the application is intended to function. My point in the "What needs to be done" section is that I believe they can still achieve the same end result in a more efficient fashion that is more respectful of shoppers' privacy.

I have just got off the phone to ParkAssist (Sydney); they were unaware and are now investigating/contacting the appropriate people within westfield.

The contact details were located using the API:)

http://120.151.59.193/v2/email-alerts

I sure hope there's authentication on the signs.json PUT update API, or Westfield's signs are probably about to go crazy. (i.e. ability to remote control the parking indicators)

Troy: Thanks for reaching out to Park Assist, I was about to contact them in addition to Westfield given the attention this is receiving. Hopefully they can plug any actual vulnerabilities quickly and address the privacy issue in due course.

"Dear Benevolent Developer, We cannot thank you enough for notifying us.

Here's what happened:

- 1. The API should never have been made accessible publicly. The regular authentication that protects API access was incorrectly disabled. We have corrected the configuration and the services are now protected again.
- 2. The code snippet should never been placed on pastie. We were sloppy in that regard.

While the information in the API is mostly used to get counts of how many spaces in the car park are currently occupied (which you can find out by going to the site and reading the digital signage), any unauthorized access to the data is an unacceptable breach.

We acknowledge that these mistakes should never have occurred and we will need to take a hard look into our security procedures to ensure this does not happen again.

If anyone would like to discuss this issue further, we welcome your comments and advice. Thanks once again. Ian"

If they're using BASIC auth without SSL, then this is still fundamentally broken; the proxying approach will reveal the auth details, so the API is still exposed.

Troy: For this app, there need not be any auth - the data is meant to be publicly queryable. Of course how they secure the existing services to the non public apps, that's another question altogether and certainly I'd like to see TLS play a role there.

Epilogue

I'd never take this approach to disclosure today, namely going public straight away without first raising concerns directly with the organisation. But hey, I was young(ish) and it was a very different time to today. That said, I only "went full disclosure" with information that couldn't be immediately exploited for nefarious purposes and there was an easy fix: just turn it off. The impact would be low, they could get it fixed then fire it back up after.

This post was the first one that got me media attention. I'd never dealt with reporters before, so this was all new to me and if I'm honest, the attention was kinda cool. I remember the story getting traction, my name appearing in the press and then... my boss coming to have a word with me. I was still at Pfizer and I'd had that boss for a decade already, but he wasn't used to one of his team suddenly appearing in the headlines. He wasn't concerned from an organisational perspective, but he warned me about how reporters would try to manipulate stories to suit their own agenda. It was an awkward discussion

that's funny to look back on now at a time when barely a day goes by without me speaking to a journalist.

The other thing that happened shortly after writing this blog post was that the General Manager for Westfield got in touch. In fact, he wrote a comment on the post which still stands there today, and he later followed up with a phone call. This could have gone downhill fast if he'd taken a dislike to how I'd written up the vulnerability. These days I joke that if something like that was to happen it would make for a good blog post and I've now got the clout to (probably) make it much harder on them than they'd make it on me, but that certainly wasn't the case in 2011.

And finally, whilst writing this epilogue I couldn't help but think how amazing it is that a decade on, I'm still writing the same sorts of blog posts about egregiously poor security flaws. Stupid little things like simple tampering of an API giving you access to crazy amounts of information. I lament that fact from the perspective of the industry just not being able to get on top of this stuff but then again, as I often joke, because of that the job security in this industry is amazing!

LESSONS IN WEBSITE SECURITY ANTIPATTERNS BY TESCO

As with so many blog posts that would follow, this one started out with one little thing antagonising me: a tweet. Mind you, that tweet came in response to my pointing out a pretty egregiously bad security flaw, namely the presence of passwords stored in plain text. But as I'd experience many times more in the years to come, once the person responsible for managing the social media channel starts tweeting about security, things tend to go downhill real fast.

So, I started picking at the site. I suspect this is what's at the very beginning of many hacks, just someone being drawn to a particular site and starting to probe away. To be clear though, this wasn't me carrying out any sort of unauthorised and nefarious activity, I was just observing the way things worked and drawing some reasonable conclusions from that. I find it just as fascinating now as I did then that perfectly innocuous use of a website can turn up so many dodgy little flaws. That was all it took to write this post, but holy cow did I later get some milage out of it!

30 JULY 2012

akedowns of companies' security postures due to crazy things their social media accounts have said have been a constant theme throughout my blogging career, and the Tesco one really got me noticed in the UK:

Update, 14 Feb 2014:

A year and a half on from writing this, Tesco has indeed suffered a serious security incident almost certainly as a result of some of the risks originally detailed here. Read more about it in <u>The Tesco hack – here's how it (probably) happened</u>.

Let me set the scene for this post by sharing <u>a simple tweet from last night</u>:



Ok then, that's about as many security misdemeanours as I reckon you can fit in 140 chars! For those wondering, yes, this is actually a verified account and it really is Tesco responding to me. I'll come back to Tesco's many interesting views on security a little later, but first, some background:

I keep a watch on mentions of my blog over on Twitter and get a lot of <u>tweets</u> <u>along these lines</u>:



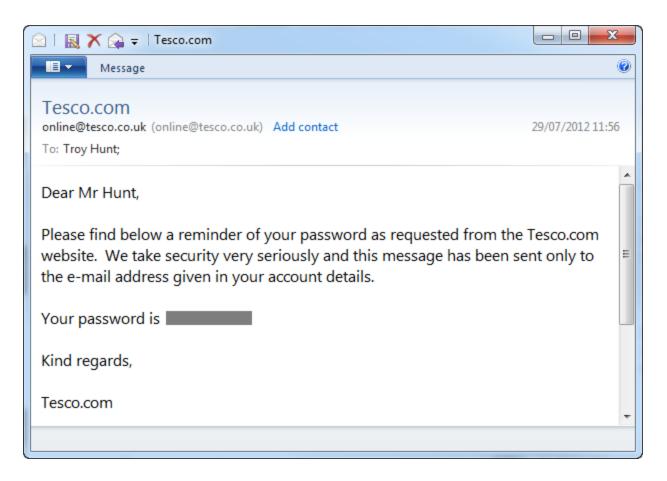
Curious, as always, I headed over to <u>tesco.com</u> to take a look. A few cursory glances around showed perhaps there was a bit of an opportunity here – an education opportunity for developers who like to learn from anti-patterns, i.e.

seeing how those who have gone before them have done it wrong. So let's take a look at the many simple security errors Tesco have delivered and see how we would approach this differently when applying basic security principles.

Oh – and for audiences outside the UK, Tesco is a major supermarket chain the likes of Coles in Australia or Costco Safeway in the US. You know, the kind of multi-billion dollar brand that *should* know how to get web security basics right, particularly when they're providing online shopping services and handling your payment info. They also provide banking and insurance services, although that's not an area I'll look at in this post.

Password storage

This was obviously the first place to start given Dan's comment. As it turns out, although I live in Sydney I actually have a Tesco account from my time back in the UK around the turn of the millennium. I wonder what my password was back then...

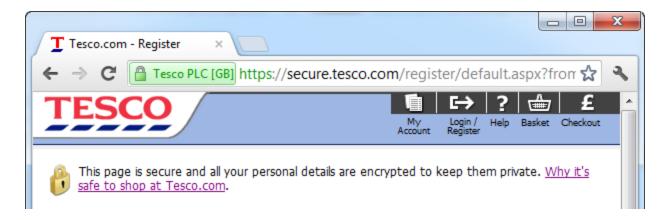


Oh dear, well Dan certainly nailed it when he mentioned no salt, clearly the passwords aren't hashed at all let alone salted. At best they're encrypted but chances are they're stored in plain text, unfortunately there's no evidence to the contrary.

Encryption is not selective - data is either important or it's not

When you decide data is worth protecting, you need to be consistent in your approach. There's no point tightly securing it in one location then having it flapping around in the breeze at another.

Apparently, encryption is important to keep your personal details private:



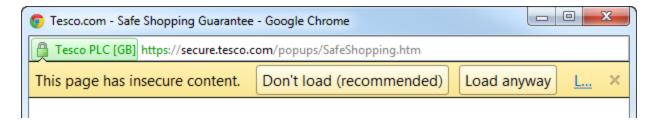
Righto, so how exactly was that password protected in email? Well of course it wasn't protected at all, it was just sent off willy nilly.

But hang on – Tesco put a big yellow padlock on the page – it *must* be secure! This is exactly the sort of thing I was talking about a year ago when I said <u>the padlock icon must die</u>. It has become nothing more than a token gesture which in no way guarantees any of your content is actually secure. It's like those people who put up signs saying "beware of the dog" when all they've got is a geriatric hamster.

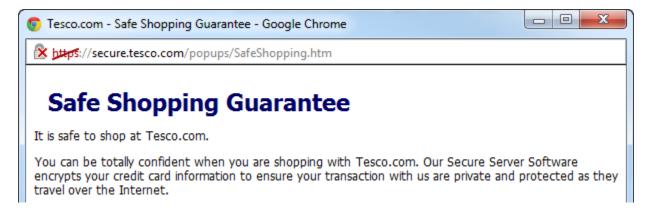
Mixed content and browser warnings

Everyone remember what mixed content is? That's when you load up a page over HTTPS – which implies a degree of security – but then it embeds resources loaded over HTTP which gives you no assurances whatsoever.

This is bad. In fact it's so bad that the browsers of today give you a very blatant warning when this is happening. Cast your mind back to the image above with the text which says "Why it's safe to shop at Tesco.com". Know what that link does? This:



Ok, so the browser is being very clear here: do not trust this page (the one about trust), in fact don't load it at all. What the heck, I like living dangerously:



When you start getting big red crosses in your browser, be afraid, be very afraid. But you probably should be even more afraid of Tesco's simple opening paragraph:



Well at least they're direct. Completely wrong (at least based on what we've seen thus far), but direct. They have a secure server (we know this because we saw their padlock) so we should all be confident that "transaction with us are private". Nice.

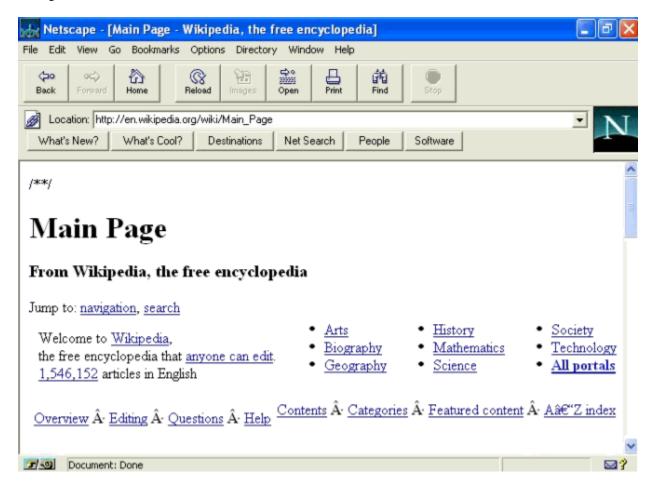
Browser version craziness

But Tesco won't just let any old browser in, oh no, you must be using something modern like version 3.0 of "explorer" or 3.02 of Netscape (this is from further

down the screen in the previous grab):

We accept orders only from Web browsers that permit communication throught Secure Socket Layers(SSL) technology; for example 3.0 versions or higher of explorer and versions 3.02 or higher of Netscape. This means you cannot inadvertently place an order through an unsecured connection.

Remember what these guys looked like? Do you even remember Netscape? Quite possibly not because there's a sizable audience out there browsing the web today who were still breastfeeding when it launched in mid 1996. Here's a recap:



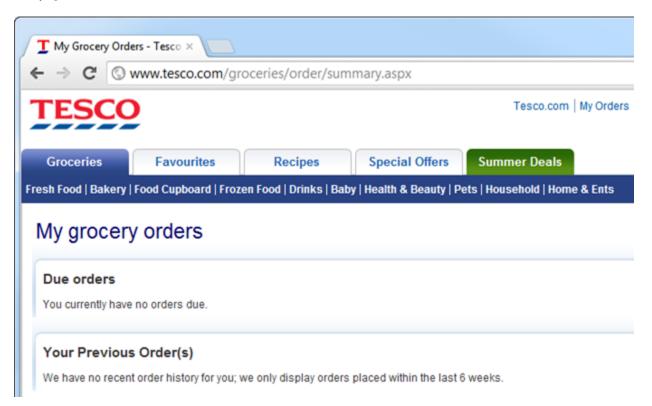
Fancy.

Why Tesco feels the need to direct users to ensure they meet a 16 year old browser version (or above, thank you) is rather odd. In fact it gives the

impression that nobody is really paying much attention to the site. Heck, it was odd when I originally created my account in 1999!

Persistent insecurity through HTTP cookies

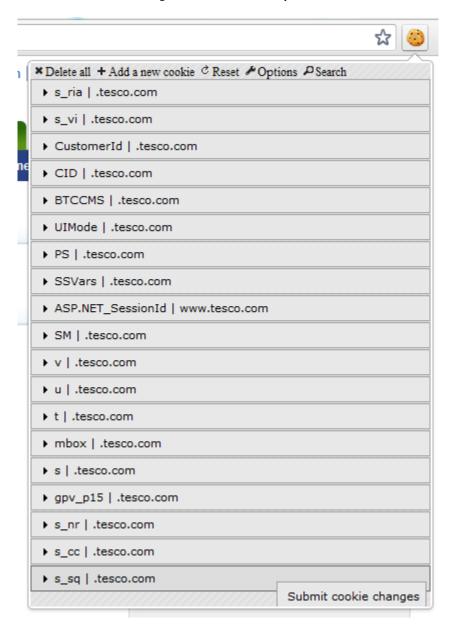
Remember how Tesco is secure? It must be because they had that padlock icon, but to their credit, they did actually provide login forms over HTTPS. But then they go and do this:



See the problem? No more HTTPS – but we're still logged in! Ah, but the password was sent over HTTPS at login so it *must* be secure, they'll say. Except, of course, <u>SSL</u> is not about encryption, or at least it's not just about encrypting login credentials.

You see, HTTP is stateless so the only (practical) way a state such as being

logged in can be persisted is by passing cookies backwards and forwards between the browser and the website. Each time the user makes a request, the browser says "Hey, I'm meant to be logged in as Troy, here's my cookie to prove it". You can see those cookies right here courtesy of <u>Edit This Cookie</u>:



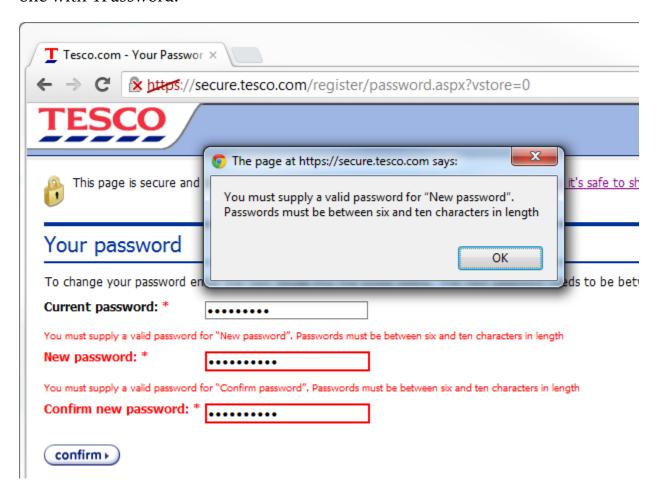
And because they're being sent over an HTTP connection, *anyone who can watch the traffic* can see those same cookies. And copy them. And hijack your session. Sound tricky? It's not, in fact I demonstrated how easy it was in <u>part 9 of my OWASP for .NET devs series on insufficient transport layer protection</u> when I

did precisely this.

Password rules

I've got to come clean – I'm a sinner. I haven't always created strong passwords. I reused them. I sometimes used the bloody dog's name, dammit! But I'm reformed and am now a fervent advocate of the mantra that <u>The only secure password is the one you can't remember</u>.

So let's jump on over to the "change password" page and generate me a strong one with 1Password:



Oh dear, 10 characters. That's it. Conventional wisdom – any wisdom – states that passwords should be long, random and unique, the more of each, the better

(and please, don't post that f***** XKCD horse battery comic as a counterargument). So what's going on at Tesco? Only 10? I'll tell you what this says to me and it goes back to the password storage point earlier on: someone's got themselves a varchar(10) under there somewhere and it's all sitting in plain text. Of course I can only speculate, but the evidence does seem to suggest this on numerous fronts.

But you know the odd thing about this? There are 10 characters in both those password fields, in fact the value is "g'Szq\5tMk". So why the problem? I've complied with the rule, haven't I?

It takes a bit of sleuthing around the site to understand what has gone wrong:

Your password:

- should be between six and ten characters in length
- can contain a mixture of letters and numbers
- treats upper-case and lower-case letters the same

Ah, letters and numbers only. Oh – and don't bother about case, that just confuses things. It's about now I'd usually get out the old soapbox and wax lyrical about the ease of brute forcing a password with these rules, but, well, you only tend to brute force a password when it's protected to begin with.

Lack of case sensitivity is also another pointer to how passwords might be stored; what are the chances that the password column in the database is simply a non-case sensitive collation and there's a query that does a direct comparison with the provided username and password? Certainly the password provided at logon is not being hashed and compared to the one in the database or that would fail the case sensitivity test (and we know this anyway because they're emailing passwords), and we also know the provided password is not being encrypted then compared or that would also fail. In fact the only real possibility that leaves any credibility whatsoever is that the stored password is being decrypted then compared to the password provided at logon using a non-case

sensitive comparer.

The other odd thing on the case front is that the password that was originally emailed to me was all uppercase. Now I know that's not how I created it, so perhaps they're just converting it to upper before storage? Or in other words, the fidelity you create in your original password entropy – even within the restrictive length and character type constraints – is lost as soon as the account is created.

But here's what I really like about that page (incidentally, it's shown during the registration process):

You can be 100% confident shopping with Tesco.com. To ensure your security online we will ask you for your password when you access your personal details or go to pay. This information will always be encrypted so that it is not possible for anyone else to access it.

Wow – 100% confident! Always encrypted! Liars.

Security misconfiguration

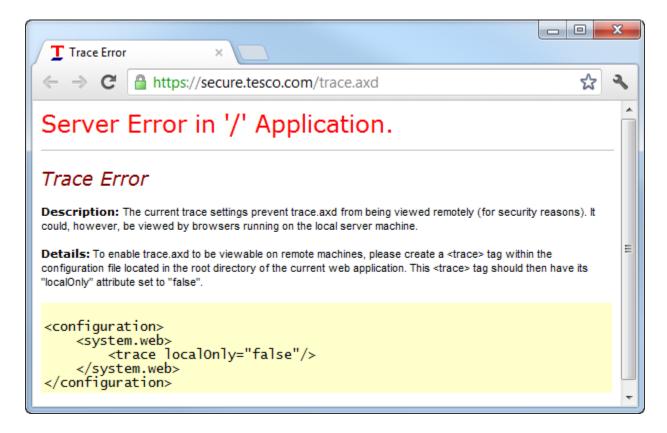
One of the things I spend *a lot* of time looking at through my work on <u>ASafaWeb</u> is security misconfiguration. This refers to those little settings available in modern day web stacks that can so easily go awry and start disclosing internal implementations which then leak information an attacker could potentially exploit.

An easy check for security misconfiguration is to see if a trace.axd handler is present. One of three things usually happens:

- 1.It's present and it's unsecured thus exposing all sorts of internal nasties.
- 2.A branded, custom error page is returned politely apologising for the inconvenience (obviously the assumption is that you've landed there by

accident)

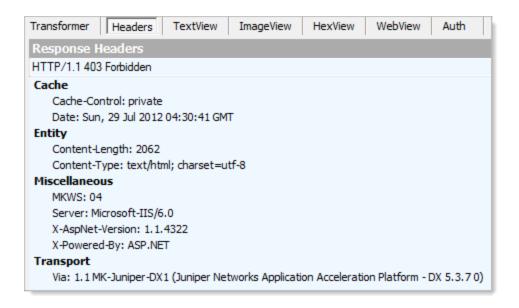
3.An internal server error is returned because although the trace handler can't be displayed, the site is not configured to show friendly error messages. It looks exactly like this:



So in short, Tesco has a case of security misconfiguration which could well leak internal implementations of the code. Nice. In fact that infamous screen above is colloquially known as a YSOD, or Yellow Screen Of Death.

Very old web server and framework

You know what the problem with today's web apps is? They talk too damn much. In fact Tesco's talks so damn much it's happy to tell you a whole lot about what's running under the covers:



What you're seeing here are the response headers returned along with the earlier YSOD. I've used <u>Fiddler</u> to inspect the headers and they reveal some telling things about the choice of technology:

- 1.They're still running on IIS 6, now a 7 year old web server and twice superseded (actually, it will be thrice superseded in the very near future when Windows Server 2012 with IIS 8 launches)
- 2.They're still running ASP.NET 1.1. This one is very surprising it's now 9 years old and superseded by .NET 2, .NET 3, .NET 3.5, .NET 4 and almost .NET 4.5.

Now none of this is to say that these were *bad* technologies in the day, they weren't, but it's like saying that your 5.25 inch floppy disk is a good thing. It had a time and a place and both of those are now gone. The security landscape has changed *significantly* since these technologies where launched and ongoing improvements in newer generations of the breed make continued progress in ensuring a more secure app by default.

Don't underestimate the value of keeping software components up to date, in fact OWASP specifically call this out in part 6 of their Top 10 Web Application

Security Risks:

Do you have a process for keeping all your software up to date? This includes the OS, Web/App Server, DBMS, applications, and all code libraries.

This is not necessarily a high-intensity exercise, once every few years you simply make sure you haven't fallen too far behind the eight ball. Certainly you don't let key software components get 9 years old and nearly 5 versions out of date.

Unconscious incompetence

Ever heard of the four stages of competence? It starts off with unconscious incompetence:

The individual does not understand or know how to do something and does not necessarily recognise the deficit.

After tweeting about security deficiencies, Tesco amply demonstrated that in terms of web security, they're firmly stuck in that first stage of competence:



I like the authoritative nature of this response and the software security prowess of the Customer Care account! The other statement that is *always* a red flag to me is "industry standard"; in my experience, this is the canonical response given by people who actually don't know the mechanics of what they're talking about

(which of course is what you'd expect from a Customer Care department). It's like "best practice" and whilst it looks good on a PowerPoint deck, once again, it is by no means evidence that you actually understand the execution of what you're talking about.

Clearly this was not one to let go without a response:



And that's when we got that zinger from earlier on:



In fact it's such a zinger that it has become rather popular:



There's probably a lesson in there somewhere about not letting your Customer

Care folks make technical statements via social media...

But this has really been the theme the whole way through; Tesco continually overstate their security prowess whilst clearly under-delivering in their execution. Yes, this is a Customer Care account but as condescending as it may sound, this really is a case of unconscious incompetence not just by a semi-automated Twitter account pulling standard lines from the book, but by the people creating the web assets. And that's inexcusable.

Lessons for developers everywhere

As I said right at the beginning, let's take a constructive view of Tesco's approach to web security and learn a few lessons along the way. I'm under no illusion that Tesco themselves will turn around and fix things in response to this post, they've known about these problems for long enough and obviously they've elected not to do anything about them.

So the lessons for developers:

- 1.Password storage should always be done using a strong hashing algorithm. IT should be one designed for password storage and also use a cryptographically random salt. It also must be a slow hashing algorithm read Our password hashing has no clothes if this is a foreign concept.
- 2.Password retrieval should never happen. Indeed it can't if you've implemented the previous step correctly. Always implement a secure password reset process. Read Everything you ever wanted to know about building a secure password reset feature for some tips on this.
- 3. Never mix HTTP content into your HTTPS pages. If HTTPS is important to you and it should be either explicitly refer to the HTTPS

protocol in your references or even easier, use protocol relative URLs. There's plenty of info in <u>OWASP Top 10 for .NET developers part 9:</u>
<u>Insufficient Transport Layer Protection</u>.

- 4. Always send authentication cookies over HTTPS. These are almost as valuable as the password itself; it gives anyone who holds them the rights to perform any tasks the user who originally authenticated to the system can. See the link in the previous point for more information.
- 5.There should never be restrictions on password entropy. Don't exclude special characters, don't chop the length at a short, arbitrary limit (if you have to, make it 100 chars or so) and definitely don't implement a system which is case-insensitive. See Who's who of bad password practices banks, airlines and more for more common mistakes.
- 6.**Ensure basic security configurations are correct**. Tracing is off, custom errors are on, a default redirect page exists, debug mode is off, etc. This is obviously for ASP.NET, but there are parallels in other web stacks. Check your .NET apps with <u>ASafaWeb</u>.

And lastly, <u>don't let your ego write checks your body can't cash</u>. Padlock GIFs and statements on web pages mean absolutely nothing if what's underneath has got holes all through it. In fact, just like that classic Twitter comment in the opening paragraph has shown, it makes things much, much worse as it demonstrates unconscious incompetence.

You know what the Tesco situation reminds me of? What I've written about above is very, very similar to what I wrote about Billabong a couple of weeks ago. Not necessarily the same vulnerabilities (although some of them are pretty close), but the rampant and readily observable failures in basic security

practices. I wrote about Billabong *after* they'd been hacked and 21,000 account details published. I concluded the post by saying:

Now I'm not saying that any of the specific vulnerabilities identified above were at the root of this breach, but what I am saying is that they indicate Billabong clearly never took security seriously. It's obvious that there's a fundamental security process missing and if such glaringly obvious vulnerabilities are present — ones that can be observed from the browser alone — what else lies within? The site was a red flag — it made it crystal clear that a bit of probing would very, very likely turn up more serious flaws.

Now think about Tesco – what can we conclude about their risk profile? Let's just say if you have a Tesco account, I'd make damn sure you haven't reused that password anywhere else.

Update, 30 July:

A reader directed me to <u>this response from Tesco's Customer Service Manager</u> a couple of years back when the issue of website security was raised with them. Here's the interesting bit:

I've had a word with my support team and asked them if they're stored with 'one way encryption' or any encryption and they say that although the information is not encrypted the level of security surrounding the password means that only the senior technical positions could access the information.

This is on Pastebin so take it with a grain of salt, but certainly the message is consistent with the level of understanding Tesco appears to have and passwords stored with no cryptographic means whatsoever would be consistent with what I've observed above.

Comments

To be fair, shopping websites aren't there to protect anyone's ID; they exist to encourage people to spend money. Doesn't matter who's money or who spends it. Trashy levels of "security" do that just fine. Mission accomplished.

Troy: Except trashy security doesn't do just fine because when they get breached and their customer database ends up on Pastebin the first thing that happens is evildoers take all those user names and passwords and start using them to hijack Twitter / Facebook / GMail accounts. Clearly the user is at fault by reusing credentials, but any organisation that accepts these credentials has a duty of care to protect them.

Great article Troy - though I have one nit-pick...

"we also know the provided password is not being encrypted then compared or that would also fail"

It is entirely possible that when the account was created, the passwords are uppercased (or lowercased) and then encrypted and stored. This would then allow the password check to do the same and compare the crypt strings.

Having said that, the evidence certainly points to plain text storage or, at very best, 2 way encryption - I just wanted to correct that one assertion.

Troy: true, although I note the new password I created was sent back to me by the reminder with the original case so it seems the old storage mechanism differs from the new.

Great writeup Troy.

Just a really minor nitpick, but "Certainly the password provided at logon is not being hashed and compared to the one in the database or that would fail the case sensitivity test" isn't necessarily the case - as you go on to say in the next paragraph "perhaps they're just converting it to upper before storage?" If they use an uppercase or lowercase version as the hash input for storage & comparison then it can still be hashed & salted.

As you do say though, that's obviously not happening here and it's quite probably a bad idea to do this at all for password entropy reasons.

The tradeoff always seems to be reasonable security vs. easily preventable user support costs. Facebook apparently made an interesting compromise of allowing you to authenticate with either yourPassword or YOURpASSWORD's case-inverted opposite or YourPassword with the first letter capitalised...

Troy: That's true, although I note that after I changed my password to a mix of upper and lower the reminder sent it to me with the case retained. It seems to only be the older accounts which upper cased the whole thing in storage.

Epilogue

Somehow, the Tesco smackdown blog post became a really big thing. It got heaps of news coverage and my name got inextricably linked with the supermarket brand, with this incident coming up in conversation quite regularly when I began frequently travelling to the UK years later. Several other blog posts followed and as of today, there are thousands of results when Googling for my name next to Tesco's.

I actually ended up meeting some Tesco security folks in London years later. It was a Pluralsight dinner at a fancy restaurant where I was speaking to a table full of executives as we wined and dined (there's a photo of this in a later blog post about my massive international speaking trip). The thing that struck me is

that despite Tesco's woefully bad attitude publicly, the guys I met at this dinner were super nice. Not just nice, but smart people who well and truly knew better than to do the sorts of things I'd called Tesco to account on. That was another recurring theme I'd find throughout my career; the technical folks inside organisations I'd publicly derided were usually really switched on. I'll talk about this more later in the post about public shaming which, as was the case with Tesco, is a practice I still find very effective at getting the bean counters out of the way and letting the smart technical people behind them do their job.

INTRODUCING "HAVE I BEEN PWNED?" – AGGREGATING ACCOUNTS ACROSS WEBSITE BREACHES

Like most developers and, I dare say, techies in general, I'd started lots of different projects over the years. I remember in about '97 sitting down and writing a website I'd called "eCars" which was intended to be an online marketplace for vehicles. A couple of years later while in London I started building a site called "Blue Shovel" which I hoped would help Aussies find tech jobs during the dot com boom in the UK. Lots of other little things followed and one of them - ASafaWeb - actually even got a bit of traction (it was a dynamic analysis tool to assess the publicly observable security configuration of ASP.NET web apps), although certainly never made any money. I had no reason to expect that HIBP would be any different to my earlier failures, but that didn't matter because I just felt like coding.

The one constant thing that prevailed across all the various pet projects I took up was this feeling of empowerment that code on the internet gave me. I still recall so vividly those very early days of the web where I was like "this is amazing, I can sit at home and build stuff that people all over the world can use!" To this day, I still feel that same sense of wonderment at the very concept, in fact now so much more than ever thanks to the magic of "the cloud". It almost feels trite writing it like this, but we live in an amazing and unprecedented time where building a service like HIBP is within easy grasp of all of us. That excites me even just writing it now, and it's that sense of excitement that led to building HIBP and writing this blog post.

04 DECEMBER 2013

often write up analyses of the passwords disclosed in website breaches. For example, there was <u>A brief Sony password analysis</u> back in mid-2011 and then our local Aussie ABC earlier this year where I talked about <u>Lousy ABC cryptography cracked in seconds as Aussie passwords are exposed</u>. I wrote a number of other pieces looking specifically at the nature of the data exposed in individual sites, but what I really found interesting was when I started comparing breaches.

In the middle of last year I wrote What do Sony and Yahoo! have in common? Passwords! and found that 59% of people with accounts in both sources used the same password. Then just last month when I wrote about "the mother of all breaches" in Adobe credentials and the serious insecurity of password hints, I found that many of the accounts from the Sony breach were also in Adobe's. In that case I explained how this put personal information at serious risk as the unencrypted password hints in Adobe's breach often had the answers in the unencrypted Sony passwords!

As I analysed various breaches I kept finding user accounts that were also disclosed in other attacks – people were having their accounts <u>pwned</u> over and over again. So I built <u>this</u>:

';--have i been pwned?

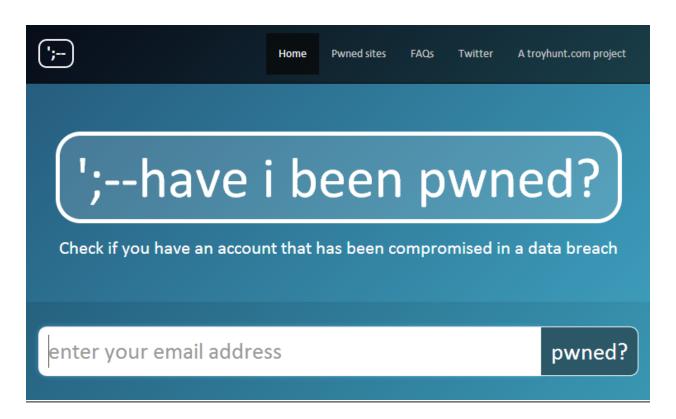
The site is now up and public at <u>haveibeenpwned.com</u> so let me share what it's all about.

About HIBP

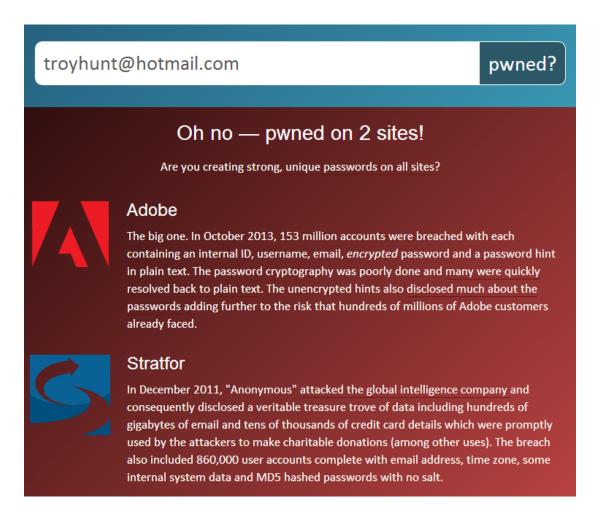
Just after the Adobe breach, a number of sites started popping up that let you search through the breach to see if your email address (and consequently your password), was leaked. For example there was this one by Ilias Ismanalijev, here's another by Lucble and even LastPass got on the bandwagon with this one. When I used the tool to check my accounts, I found both my personal and work accounts contained in the breach. I had absolutely no idea why!

The most likely answer is that I did indeed create accounts on Adobe, perhaps as far back as in the days when I was using Dreamweaver to build classic ASP whilst it was still owned by Macromedia. The point is that these accounts had been floating around for so long that by the time a breach actually occurred I had no idea that my account had been compromised because the site was simply no longer on my radar.

But of course Adobe is not the only searchable breach online, there's also one for Gawker, another for LinkedIn passwords (emails and usernames weren't disclosed) and so on and so forth. Problem is, there's not a tool to search across *multiple* breaches, at least not that I've found which is why I've built haveibeenpwned.com:



Enter your email address and go – any of the sites the address appears breached on will return a result with an overview of what happened to them. Here's an example:



As I mentioned earlier, my email address *was* in the Adobe breach. Fortunately it wasn't in any of the others so I've just added in Stratfor for illustrative purposes.

As you'll see in the footer of the site, there's rather a broad collection of accounts – over 154 million as of today – and they break down like this:

- 1.152,445,165 Adobe accounts
- 2.859,777 Stratfor accounts
- 3.532,659 Gawker accounts
- 4.453,427 Yahoo! accounts
- 5.37,103 Sony accounts

Despite the lowball reports of "only" 38 million, the Adobe dump did indeed have more than 152 million unique email addresses in it which is obviously a *staggeringly* high figure (there's some contention as to whether an "account" is only one being actively used which may account for the discrepancy). As significant as the likes of the Stratfor breach appeared at the time (and certainly it had a serious impact on them), it was a "mere" 860 thousand odd accounts and the others less again. Even so, there's a lot of commonality across the victims of the breaches.

The prevalence of multiple breaches by user

Importing the data – particularly the 153 million Adobe records – wasn't a small task, at least not to get it into the structure I wanted. I'll write more about that in the next day or two in terms of the underlying architecture, but the way I approached it was that I imported the Adobe data first and then for each subsequent breach either added new addresses or updated the existing address information about the subsequent breaches on the same account.

When I added the Stratfor breach to the existing Adobe records, 16% of the email addresses were already in the system. I moved onto Sony and 17% of them were already there. Yahoo! was 22%. Whilst not the chronological order in which the breaches occurred, what this demonstrated is that subsequent data sets showed a high correlation between new breach data and existing records in the system and that's the very reason why I created this site.

Entering the era of breach data reuse

One of the things I noticed with the Adobe breach that I haven't seen in previous cases was *other* companies notifying their users that their Adobe account had been breached. Not just one or two companies, but many of them. For example, Facebook did this and actually matched breached credentials with the ones they had on file:

Facebook users who used the same email and password combinations at both Facebook and Adobe's site are being asked to change their password and to answer some additional security questions.

I wasn't notified by Facebook (it's no surprise that I don't reuse credentials!), but I did receive a <u>notification from Evernote</u> purely because my *email address* was the same on both systems. After I wrote about the Adobe analysis, I was also contacted with requests for help in generating similar notifications for other purposes.

The point is that analysing breach data appears to be becoming mainstream. Arguably the sheer volume of the Adobe breach was the catalyst, but I do find it interesting how illegally obtained data now well and truly in the public domain is being used for constructive purposes. My hope is that HIBP can continue with that trend.

Future breaches and roadmaps

Clearly we haven't seen the last of the data breaches, of that there can be no doubt. Now that I have a platform on which to build I'll be able to rapidly integrate future breaches and make them quickly searchable by people who may have been impacted. It's a bit of an unfair game at the moment – attackers and others wishing to use data breaches for malicious purposes can very quickly obtain and analyse the data but your average consumer has no feasible way of pulling gigabytes of gzipped accounts from a torrent and discovering whether they've been compromised or not.

Of course the other thing is that I've only got five data breaches here and there are *many* more out there which I'm yet to integrate. Some of them aren't suitable (LinkedIn only contained passwords and not email addresses), but if there are others you're aware of *that are now public*, please let me know. No, don't go and breach a system in order to contribute to this project!

The ability to rapidly integrate future breaches into a common location opens up a range of other opportunities to help consumers deal with account compromises in the future. I won't go into detail now, but depending on how subsequent breaches pan out there are a number of ways HIBP can help people deal with compromised accounts *early* rather than waiting until they're potentially taken advantage of.

Other miscellaneous facts:

Passwords: I'm not storing them. Nada. Zip. I just don't need them and frankly, I don't want the responsibility either. This is all about raising awareness of the breadth of breaches.

Windows Azure: This wasn't entirely an exercise to build a service, it was a great opportunity to test out some Windows Azure features I really wanted to give a good workout. I'm *enormously* happy with the result and I'm drafting up a blog with the technical details that I'll push out shortly.

Internet Explorer 8: Yeah, sorry guys. This browser accounts for 4% of traffic to troyhunt.com, has absolutely no HTML 5 support and is well and truly into its impending crisis and ultimate obliteration. I simply didn't have the time to make things play nice in IE8 and I also didn't want to add any bloat to the site to cater for such a small, declining audience. Having said that, it will work – you can discover if an account was in a breach, it just won't be a first class experience. Or second class. Ok so it's a visual nightmare but it can still perform the key function.

No bloat: The upside to no IE8 support is that this site is very, very light! There's only just over 100kb of content downloaded over 3 requests required to make it run (another 50 odd kb and 6 requests for font-awesome and the SVG logos at the bottom of the page). I could take this down further by ditching jQuery and the full Bootstrap JS but we're talking small kb numbers that are already bundled, minified and gzipped.

Massively fast: I'll talk about this in the follow-up post about the technology, but querying those 154 million records is taking about 4ms. In fact the querying and HTTP request was going *too* fast and I had to slow things down in order to properly show the animation when you get search results.

Email validation: You can search for $\underline{a@a}$ and HIBP will give it a go. As I wrote a couple of weeks ago when I started this project, <u>email validation is a nightmare</u>. There'll be a small number of junk addresses in the system and indeed you can search for seemingly invalid addresses but better to be too liberal than too strict. The validation goes like this: got an @ symbol and stuff either side of it? Right, let me check the DB for you!

Comments

Hi, I would find it *really* helpful if I could just be notified if one of my three main e-mail addresses surfaced in another breach. Can you put in some kind of notification feature?

Then, I could simply say, "Watch me@myhost.com, someone@somewhere.org" etc. If the notification is only sent to the watched address, I don't see any security or privacy implications. But I wouldn't have to check so often if there was a new breach and if one of my addresses was contained.

Troy: I'm going to build this into the site and provide it for no charge, it's just a simple little feature.

You should add a link to "Steps to take" if an account is hacked depending on the type of breach. Pretty useful website btw!

Troy: A few people have mentioned that and it's a good point. The main thing is that depending on the breach the process is different, for example where to go to reset, what might have been breached, if 2FA is available, etc. I'll give it further thought though and see if I can improve things.

Only 8 websites? They've sent notification to customers anyway. What is the point, unless honeypot e-mails for your own scamming, Troy? For example, you do not have Dell in the list. I've got an email address which I created in 2008 especially for Dell membership. Dell been hacked numerous times, just Google 'Dell Data Breach'. Particular in 2009 my Dell dedicated email been flooded with viruses. Infact no other website sent so many viruses as Dell. I even suspected it could be done for the purpose - Dell customers often bought new Dell computers when hacked badly. Dell removed my account without notification later and I eventually abounded this e-mail address. This e-mail address checks out OK on your service. Other e-mail address used for recently hacked Adobe identified as pawned correct. But the other one I use for garbage which is I am sure compromised many times, check out OK too. So 2 out of 3 are not detected. I would not call it useful detection service in any way, bloody waste of time.

Troy: Thank you for taking the time to articulate what a bloody waste of time this is!

Regardless, let me give you a constructive response. On honeypots, there are 154M email addresses already in the system from publicly disclosed breaches. Investing the effort to build this service with the sole objective of adding comparatively small numbers to that list one by one is nonsensical. Besides, if you don't trust it then <u>follow the guidance in the FAQ</u> - don't use it!

There are countless data breaches that hit the public airwaves and countless more we never hear of. Of the ones we know about, a fraction of them result in the public disclosure of the breach data (Adobe, Snapchat) whilst many more are retained or commercialised by the attackers (Cupid Media, Target). When identifiable data is accessible, I'll add it to the site. When it's not, there's nothing to add. Simple.

If you've a vested enough interest to take the time to criticise a freely available community service both here and on Facebook, perhaps contribute something constructive and point me at the publicly accessible Dell data breach and I'll add it to the site. That, IMHO, would be a bloody good use of time.

Contact fkn Dell, Troy. You worked with Dell, not me. But something tells me you won't touch this sht.

Troy: I haven't worked with Dell and the only information I have about the breach is what you've provided here. If the data is available and you can point me to it then I'll add it. If it's not or you can't then there's nothing I can do.

Dell hack was all over the news back in 2008 and is now. And you, a security specialist, have not heard about it?!

Troy: There are an enormous number of security incidents publicly reported on a daily basis and no, I'm not across all of them, particularly not when they're six years old.

Where did you get this data from? Just curious. Do these companies actually give out the data?

Troy: Companies don't willingly give out the data, it all comes from publicly disclosed breaches. Check out the FAQs for more info.

Where can I get access to the raw data? I want to try building something like to just to understand the intricacies of building applications for cloud and using various other technologies.

Troy: You can simply Google for many of the data breaches and find them linked in via resources such as Pastebin. Have a look at the <u>Pwned websites page</u> and you'll find some pointers there.

Epilogue

This is another case of there being more to the story than what was written and whilst I've told this one publicly many times before, it deserves repeating here: Yes, I wanted to build a data breach search service, but that was only half the reason I built HIBP. The other half of the reason was to build something on Microsoft's Azure in anger, by which I mean not a "Hello World" but rather a proper web app that did something significant. My job at Pfizer was becoming well and truly stale by this time and I *really* missed building software. Plus, in my job as an "architect" (quoted because I think the title is kinda terrible), I was trying to drive the organisation towards cloud in general and "platform as a service" more specifically, and I wanted to have some actual experience with it. So, I built HIBP on an Azure App Service and put the data into their Table Storage service. That it became a popular service was entirely accidental.

A significant portion of HIBP was built on a plane to the Philippines then sitting in a hotel room there during a Pfizer trip. I just pulled my TripIt up for reference and found the trip: Nov 25-30, 2013, with HIBP then launching on Dec 4. Not a lot of lead time to build the thing yet then as now, there's really not much to it,

just a bunch of email addresses and a search box.

I had no idea how successful HIBP would become. I didn't know it would be used by hundreds of thousands of people on an average day. I didn't know it would take me to testify before Congress in the US. I certainly didn't know it would become valuable and there's a whole other story in there to be told in the future. I also didn't know how quickly things would escalate; I wrote this launch blog on Dec 4, and it was only 8 days later that I was writing about how I'd just serviced a quarter of a million visitors in only 3 days. Stuff broke. It broke but it didn't matter because hey, it was just a pet project! But it was also just the beginning with so much more to come.

NDC 2014, VIKINGS, PASSWORDS AND PINEAPPLES (AND SESSION VIDEOS)

As I began gradually speaking at more events, there was one which always stood out from the others - the "holy grail" of conferences if you like: NDC Oslo. I had this mental image of a massive event on the other side of the world, filling an entire stadium and full of the biggest names in the business. Much of this impression was built over years of listening to Richard Campbell and Carl Franklin on .NET Rocks and the tales they'd tell on the show. So, in early 2014 at the prompting of my close friend and fellow speaker, Niall Merrigan, I submitted a couple of proposed talks, neither of which I ever expected would be accepted.

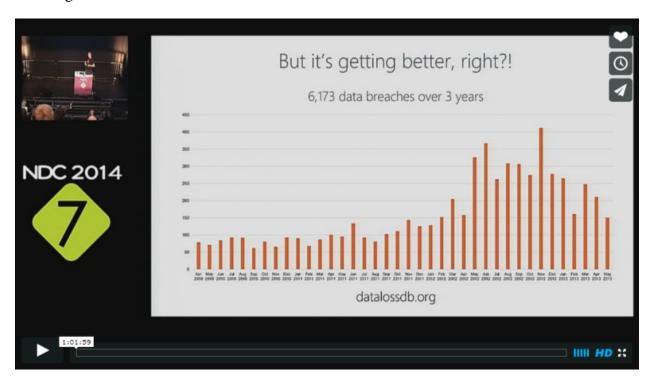
When you snare a speaking spot at an event like NDC, it's not just about getting up on stage and doing a talk; you're flown over to the other side of the world (in my case), put up in a nice hotel and welcomed into a community of (in my mind, at least) tech luminaries I'd never expected to come face to face with let alone be able to consider a peer. The whole thing was *enormously* exciting, and I invested huge amounts of effort into preparing my talks, rehearsing them ad nauseum and making sure that every little thing about what I did on stage would go perfectly. And it did.

09 JUNE 2014

ere was the original plan: propose two talks for NDC, travel over to the other side of the world and do them both then make the long trek home (each trip taking about 33 hours, thank you very much). That was pretty much how it went except that only one of the proposed talks made

the cut (I later learned that they seemed too similar which is a perfectly reasonable assessment). So I did the only sensible thing and took the very best parts out of the talk that didn't make the cut and rolled them into the one that did. And then the week before the event, they asked me to do them both. Uh...

With the originally rejected talk now cannibalised, I fell back to another recent one that had been very successful in webinar format for Pluralsight – my <u>Builders versus Breakers</u> talk. This goes through 10 online attacks, how they happened and how they could have been prevented. I find it a good talk for contextualising security risks by walking through real world attacks with real world impacts. I did this talk on the first day of the event and you can watch it now right here:



This is a good talk (at least that seems to be the consensus) and whilst sometimes the occasion calls for talking to slides as I've done here, there's also nothing like actually showing real stuff. A couple of days later I did just that with the next talk and it actually worked as a very good follow-on from the first. Wednesday was all about "here's who got pwned by SQL injection and bad crypto and insufficient SSL" then the Friday talk was "here, let me show you how

to actually exploit each of these". The second talk had a lot more humour built into it too and if the vibe from the crowd is anything to go by, I'd put this talk right up the top of the list in terms of the best ones I've delivered over the years:



This was a totally packed room – people sitting on the floor up the front, on the stairs and queuing 5 deep out the entry. Apparently the overflow room had pretty good attendance too (they have each session broadcast on eight separate screens and you listen to the one you want with headphones). Admittedly I did promote the session quite a bit in the lead-up and the promise of pwning Swedish websites was evidently alluring. That probably all contributed to having a view like this from the podium before I kicked off:



There was only really one serious glitch in the talk which was the wifi Pineapple not playing ball right at the end of the session. What should have happened was that a whole bunch of people whose devices had connected to the Pineapple would have been able to load expressen.se, attempted to login then found themselves over on my site with their (dummy) credentials on full display and the Swedish chef from the Muppets dancing around the place. I suspect the Pineapple was just overwhelmed by the number of connections; several hundred densely clustered people at a tech conference with multiple devices each will do that!

Regardless, the feedback seemed to be rather positive:



The theory is that on exiting the session you drop in either a green card for

good, yellow for indifferent or red for bad. Of those who voted (and there's always quite a few who don't, for whatever reason), at final count there were 203 green and... nothing else:)

I find that speaking is something I continue to refine after each session and by all accounts, continue to improve at. Watching the approaches of other speakers and the reactions of the audience is always interesting. The blend of humour and content, how much is ad-libbed, how much the speaker depends on static content and especially how much content there is that people can take away and actually use. For those that are interested, here are some of the talks I've seen over the last year or so that have inspired me in totally different ways:

- 1. Ben Hammersley at Web Directions 2012. This is notable for the simple fact that it remains the one technology presentation I've seen that has no slides, no demos and not a single thing on screen. In fact there was no screen and it was awesome. I really need to go back and watch it again to better understand just what it was that Ben did so well, but in an era dominated by animated GIFs, memes and live demos, to do nothing more than walk backwards and forwards on a stage for an hour in front of a captivated audience is, IMHO, a massive feat.
- 2. Erdal Ozkaya at TechEd Australia 2013. No video for this one unfortunately, but what Erdal does exceptionally well is to fill the room with an infectiously positive vibe. I heard it said in Oslo while talking with people who actually do speaker training that the audience's passion for a topic will always match that of the speaker and Erdal always does the "kid in a candy store" thing exceptionally well. He's also very engaging with the audience lots of questions, lots of direct discussion and lots of interaction. In a later session at another event, I witnessed Erdal do what to most speakers would deem unthinkable present an entire session with no visuals when the projector broke. Think about how you'd handle that, fellow speakers!
- 3. <u>Scott Hanselman at Codemania New Zealand in 2014</u>. Scott's a well-renowned speaker and deservedly so, but it's the way he goes about it that I

find most interesting. He's always extremely comfortable with the topic, that much is clear, and he injects a lot of humour into the talk that gets everyone engaged in the underlying topic. What he really does well is relays a lot of stories that illustrate his points and very rarely relies on reading words from pages so the audience is pretty much always focused on him and not the screen. What you do see on screen compliments what he's saying rather than the other way around. The subtle Microsoft-deprecating humour only helps too!

4. Nik Molnar at NDC Norway in 2014. This one from just last week was the first time I'd seen Nik of Glimpse fame talk. What I particularly liked about Nik's talk is that there was so much useful information that could be taken away from it – immediately actionable information. Of course some talks lend themselves better to this than others, but it causes you to stop and think – what are people going to actually do after seeing your talk? Are they going to say "well that was highly entertaining" yet go back to their desks and take nothing with them? Or, as in Nik's case, are they going to actually start doing things differently – better – than they did before the talk? Whilst it isn't in the recorded video, I also liked the way Nik injected humour into the blank space that normally occupies the time between the speaker being ready and the time coming to actually begin the talk (he typed out some quizzes and humorous anecdotes in Notepad).

My next talk in Melbourne next month for DDD should be better again; I'll take the bits that worked well from NDC, cut the bits that didn't and take heed of the things I learned from watching other speakers, both the good and the bad. I also had a lot of time talking with other speakers at NDC about their tips and tricks (more so than what I have at any event in the past), and there was a lot of good info in there I'll carry forward to future events.

Lastly, if you're on the fence about talking, get off your backside and just do it! Only good things happen as a result, but that's a story I'll tell at the proper time:)

Epilogue

This event ended up being a much more significant milestone in both my professional and personal lives than I ever could have imagined. Professionally, I nailed the talk. Per the blog post, I walked away with 100% positive feedback which is a rare feat, and this was my first proper international talk (I don't count New Zealand as "proper" international!) But it also happened at a time where I was absolutely fed up with my job at Pfizer. In the epilogue to the blog post about how I optimised my career, I explained what had happened with my new boss and the Codemania talk and now here I was in Oslo at the absolute peak of my burgeoning speaking career, dreading going back home to work. I sat with Niall in the hotel bar drinking whisky and pouring my heart out about how much I hated my job but couldn't see how to move beyond it. I had no idea it would be only 7 months before life changed fundamentally.

The other thing that happened on a personal level is that I met the woman I'd later marry. Charlotte was one of the conference organisers, a stunningly beautiful girl with a warm smile that welcomed everyone to the event. I recall us sitting together during the speaker cruise down the Norwegian fjords, talking about our respective homes on opposite sides of the world. We were both in serious relationships with no desire to change that status quo at the time, but we formed a friendship that saw us regularly catching up at NDC events around the world. Many years later with those relationships now behind us, we found something we never expected to find. And just to really tie this whole story together with the present day, Richard is going to marry us at our wedding in 2022. So... listening to a guy I had huge respect for but had never met helped convince me to apply to a conference where I met my future wife who I'd marry at a ceremony he officiated. Wow.

Both these stories - the professional and the personal one - have made a huge impact on my life and I wanted to make sure they were included here. Looking back, I feel like NDC Oslo 2014 was one of those sliding door moments where I could easily have done something different - for example not submitted a talk -

and my life would now be unrecognisable from what it is today.

EVERYTHING YOU NEED TO KNOW ABOUT THE SHELLSHOCK BASH BUG

I knew absolutely nothing about this topic when I wrote the blog post. I mean really absolutely nothing, but what I did know is that it was a big thing. Earlier that year we'd had the Heartbleed bug (this was in the early days of giving bugs a "brand", often complete with cool name and logo) and it was the same deal then – I knew nothing. However, I wanted to learn, and I found the best way of doing that was to write about the topic.

With both these bugs, I saw very early on that they were getting a heap of coverage in the press and were clearly big deals. But I also saw a lot of speculation and conflicting information; who was right? What really were the issues? What should we do? I spent hours and hours (probably even days), collating information about the questions people were asking and consolidating it all into a draft blog post. I researched this thing like crazy, taking on board a lot of new concepts I wasn't already familiar with and learning the vast majority of what I ended up writing on the fly. I stressed about getting stuff wrong, but I handled it by just investing more time on research until I was confident enough that I was either right, or was at least drawing very reasonable conclusions based on the information at hand.

25 SEPTEMBER 2014

Remember Heartbleed? If you believe the hype today, Shellshock is in that league and with an equally awesome name albeit bereft of a cool logo (someone in the marketing department of these vulns needs to get on that). But in all seriousness, it does have the potential to be a biggie and as I did with Heartbleed, I wanted to put together something definitive both for me

to get to grips with the situation and for others to dissect the hype from the true underlying risk.

To set the scene, let me share some content from <u>Robert Graham's blog post</u> who has been doing some excellent analysis on this. Imagine an HTTP request like this:

```
target = 0.0.0.0/0

port = 80

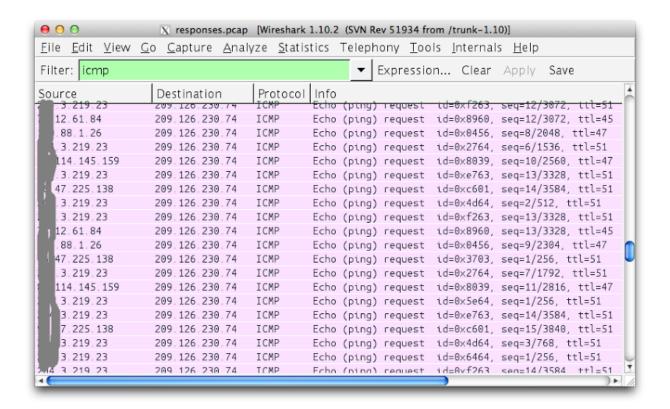
banners = true

http-user-agent = shellshock-scan (http://blog.erratasec.com/2014/09/bash-shellshock-scan-of-internet.html)

http-header = Cookie:() { :; }; ping -c 3 209.126.230.74

http-header = Host:() { :; }; ping -c 3 209.126.230.74
http-header = Referer:() { :; }; ping -c 3 209.126.230.74
```

Which, when issued against a range of vulnerable IP addresses, results in this:



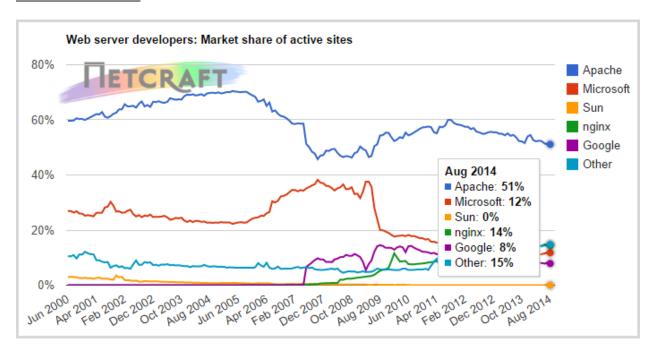
Put succinctly, Robert has just orchestrated a bunch of external machines to ping him simply by issuing a carefully crafted request over the web. What's really worrying is that he has effectively caused these machines to issue an arbitrary command (albeit a rather benign ping) and that opens up a whole world of very serious possibilities. Let me explain.

What is Bash and why do we need it?

Skip this if it's old news, but context is important for those unfamiliar with <u>Bash</u> so let's establish a baseline understanding. Bash is a *nix shell or in other words, an interpreter that allows you to orchestrate commands on Unix and Linux systems, typically by connecting over SSH or Telnet. It can also operate as a parser for CGI scripts on a web server such as we'd typically see running on Apache. It's been around since the late 80s where it evolved from earlier shell implementations (the name is derived from the <u>Bourne shell</u>) and is *enormously*

popular. There are other shells out there for Unix variants, the thing about Bash though is that it's the default shell for Linux and Mac OS X which are obviously extremely prevalent operating systems. That's a major factor in why this risk is so significant – the ubiquity of Bash – and it's <u>being described</u> as "one of the most installed utilities on any Linux system".

You can get a sense of the Bash footprint when you look at <u>the latest Netcraft</u> web server stats:



When half the net is running Apache (which is typically found on Linux), that's a significant size of a very, very large pie. That same Netcraft article is reporting that we've just passed the one billion websites mark too and whilst a heap of those are sharing the same hosts, that's still a whole lot of Bash installations. Oh – that's just web servers too, don't forget there are *a heap* of other servers running Linux and we'll come back to other devices with Bash a bit later too.

Bash can be used for a whole range of typical administrative functions, everything from configuring websites through to controlling embedded software on a device like a webcam. Naturally this is not functionality that's intended to be open to the world and *in theory*, we're talking about authenticated users executing commands they've been authorised to run. In theory.

What's the bug?

Let me start with the <u>CVE from NIST vulnerability database</u> because it gives a good sense of the severity (highlight mine):

GNU Bash through 4.3 processes trailing strings after function definitions in the values of environment variables, which allows remote attackers to execute arbitrary code via a crafted environment, as demonstrated by vectors involving the ForceCommand feature in OpenSSH sshd, the mod_cgi and mod_cgid modules in the Apache HTTP Server, scripts executed by unspecified DHCP clients, and other situations in which setting the environment occurs across a privilege boundary from Bash execution.

They go on to rate it a "10 out of 10" for severity or in other words, as bad as it gets. This is compounded by the fact that it's easy to execute the attack (access complexity is low) and perhaps most significantly, there is no authentication required when exploiting Bash via CGI scripts. The summary above is a little convoluted though so let's boil it down to the mechanics of the bug.

The risk centres around the ability to arbitrarily define environment variables within a Bash shell which specify a function definition. The trouble begins when Bash continues to process shell commands *after* the function definition resulting in what we'd classify as a "code injection attack". Let's look at Robert's example again and we'll just take this line:

```
http-header = Cookie:() { :; }; ping -c 3 209.126.230.74
```

The function definition is () { :; }; and the shell command is the ping statement and subsequent parameters. When this is processed within the context of a Bash shell, the arbitrary command is executed. In a web context, this would mean via a mechanism such as a CGI script and not necessarily as a request header either. It's worth having a read through the seclists.org advisory where they go into more detail, including stating that the path and query string

could be potential vectors for the attack.

Of course one means of mitigating this particular attack vector is simply to disable any CGI functionality that makes calls to a shell and indeed <u>some are recommending this</u>. In many cases though, that's going to be a *seriously* breaking change and at the very least, one that going to require some extensive testing to ensure it doesn't cause immediate problems in the website which in many cases, it will.

The HTTP proof above is a simple but effective one, albeit just one implementation over a common protocol. Once you start throwing in Telnet and SSH and apparently even DHCP, the scope increases dramatically so by no means are we just talking about exploiting web app servers here. (Apparently the risk is only present in SSH post-auth, but at such an early stage of the public disclosure we'll inevitably see other attack vectors emerge yet.)

What you also need to remember is that the scope of potential damage stretches well beyond pinging an arbitrary address as in Robert's example, that's simply a neat little proof that he could orchestrate a machine to issue a shell command. The question becomes this: What damage could an attacker do when they can execute a shell command of their choosing on any vulnerable machine?

What are the potential ramifications?

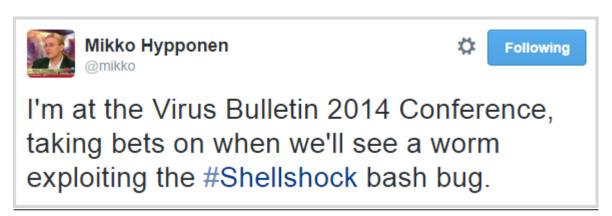
The potential is enormous – "getting shell" on a box has always been a major win for an attacker because of the control it offers them over the target environment. Access to internal data, reconfiguration of environments, publication of their own malicious code etc. It's almost limitless and it's also readily automatable. There are many, many examples of exploits out there already that could easily be fired off against a large volume of machines.

Unfortunately when it comes to arbitrary code execution in a shell on up to half

the websites on the internet, the potential is pretty broad. One of the obvious (and particularly nasty) ones is <u>dumping internal files for public retrieval</u>. Password files and configuration files with credentials are the obvious ones, but could conceivably extend to any other files on the system.

Likewise, the same approach could be applied to *write* files to the system. This is potentially the easiest website defacement vector we've ever seen, not to mention a very easy way of distributing malware

Or how about this: one word I keep seeing a lot is "worm":



When we talk about worm in a malicious computing context, we're talking about a self-replicating attack where a malicious actor creates code that is able to propagate across targets. For example, we saw a very effective implementation of this with <u>Samy's MySpace XSS Worm</u> where some carefully crafted JavaScript managed to "infect" a million victims' pages in less than a day.

The worry with Shellshock is that an attack of this nature could replicate at an alarming rate, particularly early on while the majority of machines remain at risk. In theory, this could take the form of an infected machine scanning for other targets and propagating the attack to them. This would be by no means limited to public facing machines either; get this behind the corporate firewall and the sky's the limit.

People are working on exploiting this right now. This is what makes these early days so interesting as the arms race between those scrambling to patch and those scrambling to attack heats up.

Which versions of Bash are affected?

The headlines state everything through 4.3 or in other words, about 25 years' worth of Bash versions. Given everyone keeps comparing this to Heartbleed, consider that the impacted versions of OpenSSL spanned a mere two years which is a drop in the ocean compared to Shellshock. Yes people upgrade their versions, but no they don't do it consistently and whichever way you cut it, the breadth of at-risk machines is going to be *significantly* higher with Shellshock than what it was with Heartbleed.

But the risk may well extend beyond 4.3 as well. Already we're seeing <u>reports of patches not being entirely effective</u> and given the speed with which they're being rolled out, that's not all that surprising. This is the sort of thing those impacted by it want to keep a very close eye on, not just "patch and forget".

When did we first learn of it and how long have we been at risk?

The first mention I've found on the public airwaves was this very brief summary on seclists.org which works out at about 14:00 GMT on Wednesday (about midnight this morning for those of us on the eastern end of Australia). The detail came in the advisory I mentioned earlier an hour later so getting towards mid-afternoon Wednesday in Europe or morning in the US. It's still very fresh news with all the usual press speculation and Chicken Little predications; it's too early to observe any widespread exploitation in the wild, but that could also come very soon if the risk lives up to its potential.

Scroll back beyond just what has been disclosed publicly and the bug was apparently discovered last week by <u>Stéphane Chazelas</u>, a "Unix/Linux, network and telecom specialist" bloke in the UK. Having said that, in <u>Akamai's post on</u>

the bug, they talk about it having been present for "an extended period of time" and of course vulnerable versions of Bash go back two and a half decades now. The question, as with Heartbleed, will be whether or not malicious actors were aware of this before now and indeed whether they were actively exploiting it.

Are our "things" affected?

This is where it gets interesting – we have a lot of "things" potentially running Bash. Of course when I use this term I'm referring to the "Internet of Things" (IoT) which is the increasing prevalence of whacking an IP address and a wireless adaptor into everything from our <u>cutlery</u> to our <u>door locks</u> to our <u>light globes</u>.

Many IoT devices run embedded Linux distributions with Bash. These very same devices have already been shown to demonstrate serious security vulnerabilities in other areas, for example <u>LIFX light globes just a couple of months ago were found to be leaking wifi credentials</u>. Whilst not a Bash vulnerability like Shellshock, it shows us that by connecting our things we're entering a whole new world of vulnerabilities in places that were never at risk before.

Edit: A few people have referred to the prevalence of <u>BusyBox</u> running the Ash shell on mobile devices. Devices running this don't appear to be at risk of Shellshock. The difficulty for a consumer is that they *don't know* what's running on their devices, and that includes more traditional "things" like routers as well. The long history of this bug means we've more than a couple of decades of devices out there which have gone through various evolutions of different embedded OS and we now have a *very* diverse landscape of machines and shells spanning a long period of time.

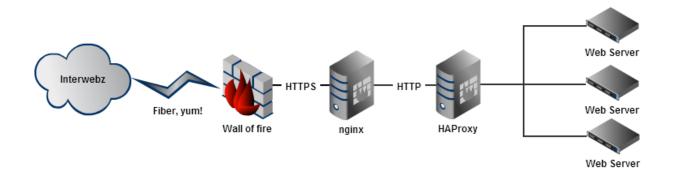
This brings with it many new challenges; for example, who is actively thinking they should regularly patch their light bulbs? Also consider the longevity of the devices this software is appearing in and whether they're actually actively maintained. In a case like the <u>vulnerable Trendnet cameras</u> from a couple of years ago, there are undoubtedly a huge number of them still sitting on the web because in terms of patching, they're pretty much a "set and forget" proposition. In fact in that case there's an entire <u>Twitter account</u> dedicated to broadcasting the images it has captured of unsuspecting owners of vulnerable versions. It's a big problem with no easy fixes and it's going to stick with us for a *very* long time.

But Bash shells are also present in many more common devices, for example our home routers which are generally internet-facing. Remember when you last patched the firmware on your router? Ok, if you're reading this then maybe you're the type of technical person who actually does patch their router, but put yourself in the shoes of Average Joe Consumer and ask yourself that again. Exactly.

All our things are on the Microsoft stack, are we at risk?

Short answer "no", long answer "yes". I'll tackle the easy one first – Bash is not found natively on Windows and whilst there are <u>Bash implementations for Windows</u>, it's certainly not common and it's not going to be found on consumer PCs. It's also not clear if products like win-bash are actually vulnerable to Shellshock in the first place.

The longer answer is that just because you operate in a predominantly Microsoft-centric environment doesn't mean that you don't have Bash running on machines servicing other discrete purposes within that environment. When I wrote about Heartbleed, I referenced Nick Craver's post on moving Stack Overflow towards SSL and referred to this diagram of their infrastructure:



There are non-Microsoft components sitting in front of their Microsoft application stack, components that the traffic needs to pass through before it hits the web servers. These are also components that may have elevated privileges behind the firewall – what's the impact if Shellshock is exploited on those? It could be significant and that's the point I'm making here; Shellshock has the potential to impact assets beyond just at-risk Bash implementations when it exists in a broader ecosystem of other machines.

I'm a system admin - what can I do?

Firstly, discovering if you're at risk is trivial as it's such an easily reproducible risk. There's a very simple test <u>The Register suggests</u> which is just running this command within your shell:

```
env X="() { :;} ; echo busted" /bin/sh -c "echo stuff" env X="() { :;} ; echo busted" 'which bash' -c "echo completed"
```

You get "busted" echo'd back out and you've successfully exploited the bug.

Of course the priority here is going to be patching at risk systems and the patch essentially boils down to ensuring no code can be executed after the end of a Bash function. Linux distros such as Red Hat are releasing guidance on patching the risk so jump on that as a matter of priority.

We'll inevitably also see definitions for intrusion detection systems too and certainly there will be common patterns to look for here. That may well prove a good immediate term implementation for many organisations, particularly where there may be onerous testing requirements before rolling out patches to at-risk systems. Qualys' are aiming to have a definition to detect the attack pretty quickly and inevitably other IDS providers are working on this around the clock as well.

Other more drastic options include replacing Bash with an alternate shell implementation or cordoning off at-risk systems, both of which could have far-reaching ramifications and are unlikely to be decisions taken lightly. But that's probably going to be the nature of this bug for many people – hard decisions that could have tangible business impact in order to avoid potentially much more significant ramifications.

The other issue which will now start to come up a lot is the question of whether Shellshock has already been exploited in an environment. This *can* be hard to determine if there's no logging of the attack vectors (there often won't be if it's passed by HTTP request header or POST body), but it's more likely to be caught than with Heartbleed when short of full on pcaps, the heartbeat payloads would not normally have been logged anywhere. But still, the most common response to "were we attacked via Shellshock" is going to be <u>this</u>:

unfortunately, this isn't "No, we have evidence that there were no compromises;" rather, "we don't have evidence that spans the lifetime of this vulnerability." We doubt many people do - and this leaves system owners in the uncomfortable position of not knowing what, if any, compromises might have happened.

Let the speculation about whether the NSA was in on this begin...

I'm a consumer - what can I do?

It depends. Shellshock affects Macs so if you're running OS X, at this stage that appears to be at risk which on the one hand is bad due to the prevalence of OS X but on the other hand will be easily (and hopefully quickly) remediated due to a pretty well-proven update mechanism (i.e. Apple can remotely push updates to the machine).

If you're on a Mac, the risk is easily tested for as described in <u>this Stack</u> Exchange answer:



It's an easy test, although I doubt the average Mac user is going to feel comfortable stepping through the suggested fix which involves recompiling Bash.

The bigger worry is the devices with no easy patching path, for example your router. Short of checking in with the manufacturer's website for updated firmware, this is going to be a really hard nut to crack. Often routers provided by ISPs are locked down so that consumers aren't randomly changing either config or firmware and there's not always a remote upgrade path they can trigger either. Combine that with the massive array of devices and ages that are out there and

this could be particularly tricky. Of course it's also not the sort of thing your average consumer is going to be comfortable doing themselves either.

In short, the advice to consumers is this: watch for security updates, particularly on OS X. Also keep an eye on any advice you may get from your ISP or other providers of devices you have that run embedded software. **Do be cautious of emails requesting information or instructing you to run software** – events like this are often followed by phishing attacks that capitalise on consumers' fears. Hoaxes presently have people putting their iPhones in the microwave so don't for a moment think that they won't run a random piece of software sent to them via email as a "fix" for Shellshock!

Summary

In all likelihood, we haven't even begun the fathom the breadth of this vulnerability. Of course there are a lot of comparisons being made to Heartbleed and there are a number of things we learned from that exercise. One is that it took a bit of time to sink in as we realised the extent to which we were dependent on OpenSSL. The other is that it had a very long tail – months after it hit there were still hundreds of thousands of known hosts left vulnerable.

But in one way, the Heartbleed comparison isn't fair – this is potentially far worse. Heartbleed allowed remote access to small amount of data in the memory of affected machines. Shellshock is enabling remote code injection of arbitrary commands pre-auth which is potentially *far* more dire. In that regard, I have to agree with Robert:





This 'bash' bug is probably a bigger deal than Heartbleed, btw.

It's very, very early days yet – only half a day since it first hit the airwaves at the time of writing – and I suspect that so far we're only scratching the surface of what is yet to come.

Comments

Troy, you only provided the first of two tests the Register suggested, and unfortunately, it's quite possible to not see "busted" but still be vulnerable. The test you provided:

```
env X="() { :;} ; echo busted" /bin/sh -c "echo stuff"
```

does *not* show busted on my system, but this test:

```
env x='() { :;}; echo vulnerable' bash -c 'echo hello'
```

does show vulnerable. I'm concerned that many of your readers may think they're ok by running the test you provided.

This is most likely because your default shell (which is often symlinked to `sh`) is *not* bash.

More likely this is because he already has installed the early but incomplete bash patch which closes form one but remains vulnerable to form 2 https://news.ycombinator.co...

Yes, I did just verify that /bin/sh was linked to dash on one of my older, unpatched, Ubuntu servers. I only run Ubuntu servers, so maybe they're *all* using dash instead of bash for / bin/sh. I'm still not 100% sure if that closes the door on the vulnerability though - I wouldn't be too surprised if some app was invoking bash directly.

Indeed, the default shell on Ubuntu (and I believe other Debian-based distros) is dash. otoh, I'm still unsure dash doesn't have the same vulnerability.

I've tried dash, ash (busybox), zsh, pdksh, and ksh. None of them have this vulnerability. You should all note that this is *not* part of the POSIX shell spec, it's a bash "feature".

The default shell for users on Debian-based distros including Ubuntu is generally bash. But / bin/sh is linked to Dash, and most system shell scripts use that.

I would strongly suggest you change your test comment from using /bin/sh to using /bin/bash. Just because the default shell is not bash, does not mean that scripts won't use it explicitly. You shouldn't assume you're safe just 'cause there's no /bin/sh -> /bin/bash symlink. And, in the Windows world, you might be surprised. There's cygwin, win-bash, MKS Toolkit... MKS Toolkit is explicitly installed by, say, Oracle iStore. While it is unlikely a windows machine is vulnerable, it certainly makes it harder to catch, since the executable is less likely to be in /bin

Naked Guy: OK, stupid question time: is this only exploitable if you're running something like Apache's CGI module, where an untrusted person can pass data to an environmental variable that's seen by bash?

I see demonstrations of how to see if your version of bash is vulnerable, and how to exploit

this via CGI, but it's not clear to me if there's any other way this could be abused. If I'm running a straight up Rails app behind a bare-bones nginx, for example, is there any danger?

I realize that as part of a layered security strategy, I should upgrade/patch bash reasonably soon, but in terms of "Am I naked to the world right this instant?", I can't tell that I am. Am I mistaken?

Paul: Yes, and on OS X the root user isn't even a valid user by default, AND apache runs as "httpd" user so I just can't see how an OS X consumer system can possibly be exploited. And does anybody use "web sharing" on a mac?

RT: Exactly. On many (most?) Unix-like boxes the most vulnerable services run under their own user account, not root, as a security precaution, so the system being compromised is not an issue.

Nevertheless, that doesn't mean that the hacker could not gain access to sensitive data that the "httpd" account has access to. This issue is a concern, but it's not a HUGE concern.

Well, if your /bin/sh is bash, and you're running a DHCP client, most likely there are some shell programs that are being run when certain DHCP admin packets come in (take a look at some of the scripts in /etc/dhcp or /etc/dhcp3). Guess what -- they're run as root. All you need is a rogue DHCP server on a public wifi.

Well that's certainly one argument for not using bash for /bin/sh. In that sense, this is a victory of the Free Software community over commercial software. Ubuntu and other distributions based on Debian wisely do not use bash as /bin/sh, whereas commercial Unix-like systems -- such as RedHat and Mac OSX -- apparently do.

Philosophically speaking, however, if a DHCP client is passing stuff that it receives to a shell script running as root without doing a sanity check on it first, then the real security hole is in the DHCP client, not the shell. I haven't had the time to look at the source code of various

DHCP clients to verify that this is the case (I'm not convinced from the documentation that it is), but if it is, then these clients should be fixed and updated, just as bash should be.

"In that sense, this is a victory of the Free Software community over commercial software. Ubuntu and other distributions based on Debian wisely do not use bash as /bin/sh, whereas commercial Unix-like systems -- such as RedHat and Mac OSX -- apparently do."

I appreciate the irony -- FSF-developed bash is the default shell in the more commercial versions of linux. But I think it's more of a victory of simplicity over a complete lack thereof: bash was bloated 20 years ago, now it's so big with so many features nobody can find all the possible traps. I was uncomfortable using the busybox shell for CGI in a router; bash would be out of the question!

Take a look through the bash variables in the reference manual: http://www.gnu.org/software... want to rename a command? Try BASH_ALIASES or BASH_CMDS. Try uploading a file to a server (if you can figure out the random name some of them give the file) and then set BASH_ENV to it. We all expect that services like DHCPC and others are smart enough to prevent arbitrary injection into these variables, but who knows what the next "cool" feature of bash will be. Bash should be used as an interactive command interpreter, not a script engine. Remember, the guys developing it are the same ones who added a lisp interpreter to a text editor.

No arguments here about the bloat in bash. I use it as my primary CLI mostly because of its ubiquity, but I could do well without over half of its "features."

It's a shame that the commercial vendors didn't learn from the Bourne/Korn shell paradigm of the commercial Unices of the eighties. The fact that Debian-based systems do not use bash as the default /bin/sh is not an accident. When I was a Debian developer, it was bug that MUST be fixed if a shell script depended on bash's extra features without explicitly calling / bin/bash. And now we see the wisdom of this decision.

By the way, Emacs is not a "text editor"; it's a way of life -- more like a philosophy or religion than a mere application. But whatever it is, it's far better and more useful than the various

godawful IDE's that less enlightened programmers cling to.

Epilogue

The blog post ended up getting a heap of views because it was comprehensive at a time when there was a vacuum of good information. As of today, there's 257 comments on the post which is a pretty good empirical measure of how much interest there was in it. Based on the success of the "everything you need to know about..." approach, I went on to write more blog posts along these lines including one about the POODLE bug the next month (this one was the nail in the coffin for SSL 3), the Apple versus the FBI case in 2016 (this was in the wake of the San Bernardino shooting where the feds tried to force Apple to unlock the suspect's iPhone), and then the WannaCry ransomware the year after.

The Shellshock post also resulted in a Pluralsight course with them reaching out after I'd written it and asking me to put something together for their customers. I teamed up with my friend Jim Manico to create a course that's still up there today. This was another example of creating something for no reason other than to put content out to the community for free, but it then accidentally turning into something of commercial value. I don't think I actually made much money out of that course (it was pretty niche), but that was a nice unintended consequence.

</PFIZER>

I'd always been really cautious to separate my online identity from my day job. Pfizer was (and I suspect still is) a very traditional organisation with traditional views and frankly, I didn't want to cause trouble for myself. I never wanted to end up getting called into the boss's office and asked to explain why I'd just written something publicly that might either reflect badly on them or be perceived as somehow disclosing a company secret. The "perceived" bit is what worried me most because whilst there was no secret sauce in the technical stuff I was doing there, I could imagine a pointy-haired manager taking umbrage with me publishing some code to my blog because, well, I don't even know but I could just see that happening. Even in writing this blog post, I was cautious with what I said as my redundancy payout was yet to land and the last thing I wanted to do was jeopardise that. But this was such a pivotal turning point in my career that I had to write something, I just watered it down to be as factual and non-offensive as possible. So, this blog post was tailor made to speak about something I'd never mentioned before in a way that wouldn't cause me problems. If you want to know what really happened, read the epilogue...

15 APRIL 2015

oday marks two important milestones for me – it's the first time I've ever mentioned <u>Pfizer</u> on this blog and after 14 years, it's my last day working for them. Both those milestones are significant and in their own ways, mark a pivotal point in my career. For those that are interested, I'd like to tell you what I've been doing in recent years and give a hint of what will come next.



"Architect"

There's this odd thing that tends to happen in many peoples' careers and I suggest it's particularly prevalent in technology: you get really, really good at something and then it hits you – you have to stop it. Well actually, you could continue doing it, but not if you want to "progress" against traditional measures such as seniority and income. That's an unfortunate aspect of our technology field in general and that was where I found things heading as I become an "architect".

Now architect, to my mind, is always a bit of a funny word in that it means very different things to different people. Does it mean you sit there drawing UML diagrams all day? Or designing what projects should go into a Visual Studio solution? Or just corralling developers into some sense of a common direction? In many cases, it means <u>architecture astronauts</u> which in my experience, are a particularly dangerous lot and certainly not a direction I was ever going to be heading.

For me, as my role became more about architecture and less about development it increasingly focused on the longer term goals of how we implemented technology and how we set a common vision for the organisation rather than just arriving at our future destinations by default. In a place worth a couple of hundred billion dollars, that can pose some rather interesting challenges. That my tenure covered such fundamental shifts in technology as the emergence of social media, mobile and cloud made things particularly interesting.

As of today though, all that comes to an end.

The exit

For most of the last decade, my role focused on Asia and other "emerging markets" across the globe. In many ways it was the best of both worlds – I got to live in what I (and many others) consider the best place on earth down here in Australia but I also got to work with countries that were growing rapidly and presenting some really interesting technology challenges as they did. It meant that in the one day I might be working out how we handled communications with customers in India where internet connectivity is low but everyone is on SMS then ensuring we did the right thing in China when it came to the government's expectations of where web assets were allowed to be hosted. The local technology idiosyncrasies across the region made it an exceptionally diverse role; Indonesia almost exclusively uses Firefox (although Chrome is making headway and IE, well...). The Philippines has a massive affinity with PHP. Pakistan has serious internet connectivity problems. Everywhere has their own unique little differences when it comes to technology.

But there were two immutable facts about my role at Pfizer and that of others in Australia who are moving on with me: I was in an expensive country which commands high wages on the global scale, and I was looking after a region which whilst growing quickly, remains cheap. There's also the fact that Australia is a hell of a long way away from *everywhere* – you can't get to Asia in less than 8 hours and that makes for very expensive journeys when they're needed, not just in terms of dollars but in time spent watching movies on the plane and not actually being productive. Add to that the fact that in the company of Asia,

Australia is clearly not where the growth is across *most* industries, and the writing was on the wall – the role was made redundant.

There are different terms for redundancy in different parts of the world, but in short it's the employment equivalent of "It's not you, it's me, but let's still be friends". It's funny watching the gears in peoples' heads work when you tell them this:

"Whoa, is that a bad thing? Should I tell him I'm sorry for him? Or wait – could that actually be a good thing? Should I say congrats..."

In my case, it's a very good thing! To their credit, Pfizer's redundancy package is *extremely* favourable to the effect that if I so desired, I could now go surfing for a very, *very* long time. In fact it's so favourable that just the potential of realising this eventuality one day kept me from doing things I really wanted to much earlier lest I miss the chance to exit in this fashion and create a whole bunch of extra choices for me. Particularly last year as things started really firing on other fronts I'll talk about in a later post, the temptation was strong and I came close... but stayed. That turned out to be a massively fortuitous decision.

So it's a big win-win; Pfizer gets to focus their energy and their dollars in the areas that make the most sense for them and the redundancy helps me to focus my energy on the things that really excite me and offer a whole new level of opportunity. I'm sure there'll be those that lament the decision (I *hope* there are otherwise I really haven't made enough of an impact) but the path forward is now clear and I couldn't be happier.

Things I will not miss...

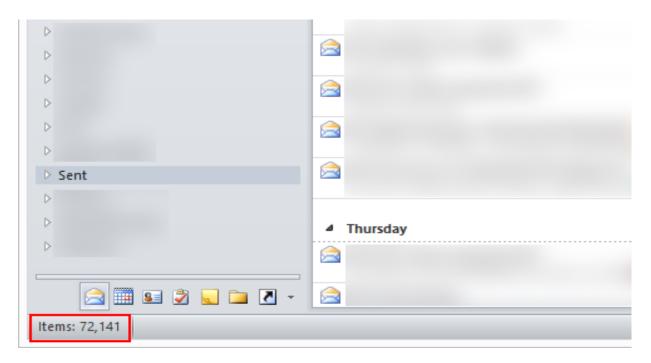
There are certainly things I can happily talk about not missing when it comes to my "old" way of working in a large corporate.

Driving to the office – I actually don't mind driving, but routine is what gets me. I

worked a lot from home in mornings and evenings and usually travelled the 30 min journey after and before peak hour, but by the time I made myself presentable, shut down what I was doing then fired back up again in the office and repeated the whole thing in reverse at the end of the day, there was a one and a half to two hours burned. Well not entirely – listening to podcasts on the drive has been enormously beneficial – but it's not exactly productive time.

Conference calls. If you don't understand why, check out <u>conferencecall.biz</u>, I'll wait. See what I mean? I'd frequently spend a couple of hours a day on the phone and I'd *always* be dealing with bad lines, incorrect conference codes, dogs, roosters (yes, you know who you are!) and other things that kept me from *actually producing things*.

Email. Now of course I'm never getting away from email entirely, but there was rather a lot of it:



That was as of last weekend and between conference calls and emails, a huge amount of my time went on *talking about doing stuff* as opposed to *actually doing it*. In fact a few years back I ran <u>RescueTime</u> for a while and found I was in Outlook for about 2.5 hours a day. That's just in Outlook and it doesn't include the time

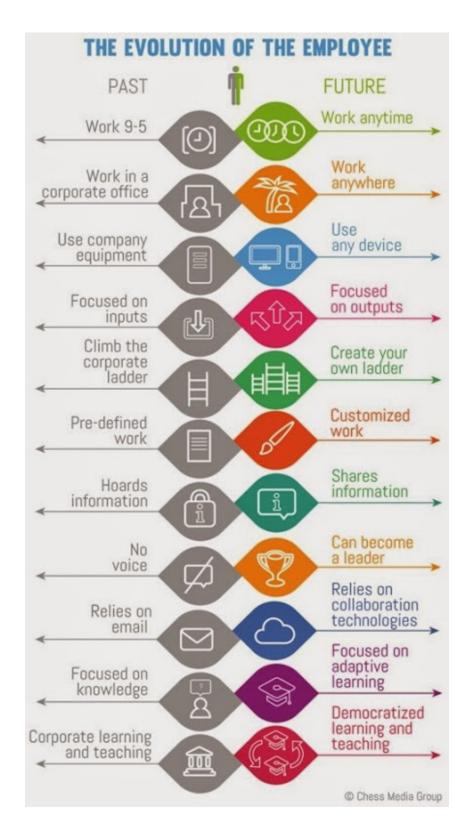
spent in other apps in order to deal with the stuff that came in via Outlook! Maybe that's what people do when they "progress", but it's never quite the same as actually creating something that's a tangible outcome.

What all of these things have in common is that whilst they were required to keep the corporate technology wheels turning, they meant that what I actually *produced* was enormously incommensurate to the *effort* I put in. That can be a bitter pill to swallow for results-orientated people, which brings me to the next phase...

The future

I've known about this exit for a few months now and what's surprised me is just how many people I interact with online didn't know I had a "normal" job. All that blogging and speaking and other things I've produced have happened late at night, on weekends, on *vacation* (most of my speaking has been done out of my annual leave) and basically whenever I could find spare moments amongst the commitments of a demanding job. I'll write more at a later date about how I multi-tasked myself to do all this, suffice to say it has been *a lot* of hard work.

So where to next? Let me talk about *how* I'll be working next and I really love this graphic from <u>Jacob Morgan's "The Future of Work"</u>. I find myself standing where the half-grey, half-green bloke is right now (although in fairness, Pfizer did some of the grey bits better than the image suggests):



There are a few things in particular there that really resonate with me and will drive my future direction. Working any time and anywhere, for example. The

reality is that I'll always "work" much more than your classic 40 hours a week anyway because I love most of what I do in the technology world and frankly the line is often a bit grey between work and play. I'd like to go for a surf during the day and take the kids to school and work when it suits my other activities. Sometimes that will be from the base of a mountain after a day on the snow.

What really clicks with me though is the ability to "create your own ladder and become a leader". The description above is pretty apt on both the past side of things and the future. In fact the reality of it is that by you being here reading this and possibly having been following me via my public presence, you've seen this happen anyway. Becoming a public identity and carving out my own niche has been enormously empowering and that is *very* firmly the direction I'll continue to head in now. My success will be directly tied to my results, not to which row I sit in on an HR spreadsheet. Many people are very comfortable with the certainty of a corporate career and that's just fine, but it was never really my cup of tea.

Identity and profile

I read back through my first public blog post ever the other day – <u>Why online</u> <u>identities are smart career moves</u> posted in 2009 – and it really hit me just how much of a turning point that was. It was a recognition that an independent identity beyond the one you have in your place of work is valuable and it also set me on a course that has given me fantastic opportunities in my post-Pfizer world. I'm going to wait until a later blog post to talk *specifically* about what I'm doing next, but let me finish this post right back where I began nearly six years ago:

The thing is though, building an online profile is not an overnight process and I don't know if I'm still going to be as enamoured with my job (or my employer as enamoured with me!) in two years, five years, ten years;

whatever! It takes a lot of time to build a public identity and waiting until you actually need one is just not going to work.

The next phase is well and truly taking shape and it's all because of this. Profile, relationships and reputation have lead to fantastic opportunities that make this a very easy, very *rewarding* transition.

I'm enormously excited by what's happening next. The experiences at Pfizer have shaped that future and exiting in this fashion is the best result I could ever have hoped for. I'll follow up with another post as the dust settles on this outgoing phase of my life, I'm *enormously* excited about what's coming next!

Comments

Congratulations for the future. It was your blog that got me interested in developer security, and as a result I'm managing to have a reasonable effect on driving our medium size enterprise forward. It's very easy to get stuck in a rut and stagnate. Your enthusiasm and easy to digest style has really motivated me to push myself forward, perhaps to do my own thing in the evenings. My commute is 3 hours (or 4 hours by bicycle) so time is tight. I can't wait to see how you juggle your time with kids, blogs, personal projects, speaking, etc, as this is what I find hard.

Troy: Thanks Dash, it sounds like you know how tricky it can be trying to juggle all these things. It takes compromise and sacrifices on many fronts. I have a post in draft that talks about how I did what I've done so hopefully that will help put things in perspective a little more, suffice to say that it has involved many years of simply putting in long hours and hard work. Of course there's more to it than that as well, let me write it up over the coming weeks.

Epilogue

I hated my job at Pfizer by the time I departed, and nothing better illustrates why than the experience I'd had only one year before writing this post. I was in the very early stages of carving out a speaking career and I'd been invited to present the locknote at Codemania in New Zealand, my first ever international appearance. I discussed the event with my boss and made him aware of the speaking engagement which was to occur last thing on a Friday, adding that I'd then spend the weekend and the following Monday taking a holiday in Auckland. This was all planned well in advance, far enough out in fact that by the time the event drew near, I'd changed bosses. For the first time in 13 years, I found myself reporting to someone else.

Gerard was based in the Philippines which meant that not only was I now reporting offshore, but I was also reporting into a totally different culture with very different norms around things like organisational hierarchy. I made him aware of my impending travel plans then, shortly after, he asked me to come up to Manilla for a meeting which conflicted with Codemania. The discussion – all in writing via email – went like this:

Gerard: "I'd like you to attend the meeting in Manilla"

Me: "As you know, I'll be in NZ so unfortunately that's not feasible"

Gerard: "I'd really like you to attend the meeting in Manilla"

Me: "I've made a commitment to speak, they're literally promoting the event with me as one of the headline talks"

Gerard: "The only priority is the company"

That last line is absolutely verbatim what was said. No joke. This wasn't going to go well as there was no way I was missing that talk. But culturally, Gerard was used to telling someone beneath his seniority level to do something and they'd just do it. It's a very hierarchical sort of place, the Philippines, and I vividly remember my trips there where I'd wander around the Pfizer office and people

would refer to me as "Sir Troy". People who perceived themselves as lower on the corporate rung than me, that is, and I suspect Gerard thought he'd get the same sort of treatment. From an Aussie. I pushed back hard, telling him that his views were inconsistent with corporate values (the whole "work / life balance" thing was becoming big then) and that there was no way I was missing Codemania. Eventually, we compromised with me still going to the Philippines but leaving a day earlier and arriving in Auckland towards the end of the conference. I had to fly via Singapore where the plane was delayed, arriving in NZ late with me then rushing to the venue, arriving last thing in the day and barely making it in time to deliver the closing talk. I was furious.

First thing I did when I got back into the office was march up to HR and get the forms required to quit. I was done. There were a whole bunch of other reasons that spread over multiple years, but the Codemania situation was the final straw. Then I saw the redundancy provisions; if they asked me to leave instead of me leaving on my own, they'd have to pay me out big time. So, I sat on the paperwork and bided my time, something that became easier to do with Gerard no longer talking to me (after 13 years there, I worked pretty self-sufficiently). Only 8 months later the redundancy came, paying me out the equivalent of almost 2 years salary and giving me the freedom to pursue my independence. Had I not made the sacrifice to stick with something I hated in the short term for the betterment of the long term, who knows where I'd be today.

I don't regret the time with Pfizer because without it, I wouldn't have had the push to build out my independent life. If I'd gone in and been happy every day, I probably would have poured spare energy into my job rather than all the other things I ended up doing that I eventually made a livelihood out of. It sucked working at Pfizer, and I couldn't have wished for a better outcome \bigcirc

HOW I OPTIMIZED MY LIFE TO MAKE MY JOB REDUNDANT

This was a very personal blog post to write. It's probably one of the first where I really started opening up about the more private side of my increasingly public life. In mid-2015, I had this newfound independence where I was finally free of the corporate shackles and actually doing, well, whatever the hell I wanted to! It was a very exciting time, too, because I really didn't know where it was all going. I had a good runway of cash courtesy of the Pfizer redundancy and the Pluralsight royalties were well and truly paying out by now, but beyond that I had no idea where I was going either professionally or financially, I was just doing what felt right at the time.

But I felt in control. I was really pouring my heart into everything I did, pumping out a lot of hours blogging, speaking, working on Have I Been Pwned and just generally saying "yes" to everything that was offered. I was biting off more than I could chew then chewing like crazy. I wrote this blog post in part to get things straight in my own mind about how I was managing my time and doing so much, and in part to encourage others to do the same. I added little bolded tips to the end of each section, outlining the key behaviours that were working for me with the hope that others could put them to good use too.

17 JULY 2015

f you're a regular reader, you may have noticed a rather major job change on my behalf recently. The day to day office grind has gone and corporate life is now well and truly behind me, where it will firmly stay. One of the things that amazed me most when I finally wrote about this is how surprised so many

people were that I actually had a normal day job:

Can't believe <u>@troyhunt</u> had another job as well! <u>#inspiration https://t.co/</u>
<u>918HOFSGLA</u>

-- Conrad Jackson (@conradj) April 15, 2015

<u>@troyhunt</u> Put me in the category of "didn't know you had a day-job apart from the blog/speaking/etc". Wow.

-- Tommy Williams $\widehat{\mathbb{V}}$ (@theRealDevgeeks) April 15, 2015

@troyhunt no one can imagine that much awesome work on the side. Hah.

-- Tommy Williams V (@theRealDevgeeks) April 15, 2015

<u>@troyhunt</u> Congrats! Count me among those who were clueless that you had a "normal" job. Are there 32 hours in a day in Australia? :-)

-- Sandra Vigil (@SJVigilant) April 14, 2015

<u>@troyhunt</u> Best of luck Troy. I'm one of those surprised to discover you had a full time job.

-- Jenny Luca (@jennyluca) April 14, 2015

Wow. Who knew <u>@troyhunt</u> had a day job! I thought he was busy enough for a full-time blogger

-- David Wengier (@ch00k) April 15, 2015

<u>@troyhunt</u> <u>@chrismckee</u> Stunned you had "normal" job as well as the fantastic blogging, training & research. love to chat re: API Security ops

-- Mark O'Neill (@TheMarkONeill) April 15, 2015

I want to write about how I did this. This is not just about how I managed my time to do so much, but how it enabled me to get to the point where I could no longer justify working in the corporate world. I made my job redundant long before Pfizer did and by good fortune (and admittedly some good management as well), I exited in the best way possible. Here's how I managed to do so much in the lead-up to that so that when the time finally came, I was in better shape than I could have ever imagined.

I multithread and task-switch frequently

I've obviously had a lot of parallel stuff on the go at once; multiple blog posts, speaking events, community interactions, <u>HIBP</u> then naturally all those Pluralsight courses *and* a full time job. I regularly switch between all of them which means I might be bang in the middle of doing something in HIBP then have a great idea for a blog post so I'll go and churn out a para or two there then jump back. I get an itch that I want to scratch but am happy then flicking back over to the other context maybe 15 minutes later.

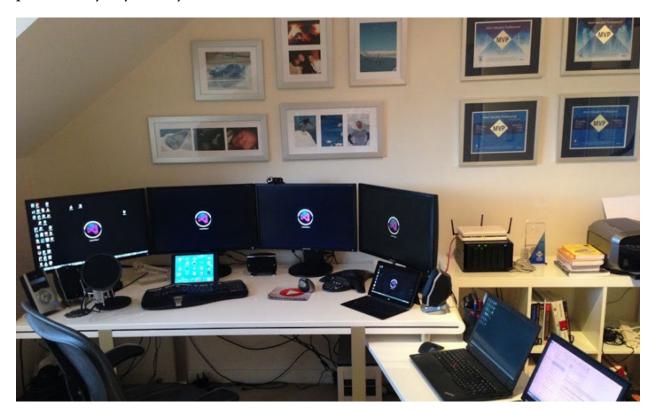
I'll also tackle multiple things at once. That might mean being on a conference call and also responding to emails (I'm happy listening and writing at the same time) or doing the social media bit while waiting for an HIBP deployment to run. Rapid context switching and multitasking have been extremely useful in many scenarios. I'm fully conscious that this way of working isn't for everyone, but I find it very effective for me.

I would say however, that there are definitely times where 100% of focus needs to be directed at one thing. Recording courses, for example, means that all the buzzy things get turned off and all alerty things closed. That's not just for the sake of clean audio, but it's for my own concentration. The trick I find is knowing when I can juggle multiple things versus when I need to focus and that's determined both by the task and my mood. Mood in particular really determines what I'm working on and that's a benefit of having so many simultaneous interests; I can be creative, analytical, conversational or reflective depending on what I feel like at the time.

Tip: Diversify your interests and the mediums you work across such that you always feel like doing *something* of value.

I tailor my work environments to maximise productivity

Being as productive as possible in the time I have available is enormously important. Many years ago now I wrote about <u>Building the ultimate virtual office</u> and I talked about things like having a good chair; a <u>Herman Miller Aeron</u> costs some cash but it'll last a couple of decades and I spend a huge amount of time sitting in it so on a dollars-per-hour-of-arse-on-seat scale, it's about the best ever use of cash. Same again with multi-monitors and these are dirt cheap these days, particularly if you buy standard DPI ones like I have.



I also bought a fast Lenovo W540 last year because maximising productivity

while I'm travelling is massively important. It's not just that, I often lay on the couch and punch out emails and as nice as machines like the Surface Pro 3 are (I actually bought my wife one), they just don't cut it for working on your lap and they suffer from a distinct absence of screen real estate for the things I like to do.

Tip: Make the places you spend your time working as effective as they can possibly be; *invest* in this.

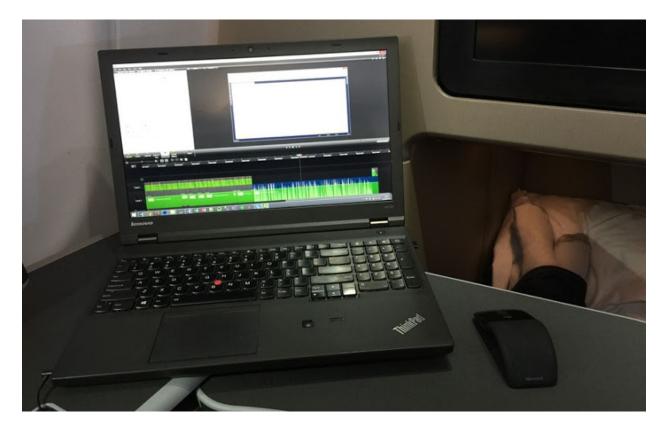
I have a really good sense of what my time is worth (and I'm willing to pay for it)

Time is money, right? Ok, that gets cliché sometimes but when it comes to figuring out where to invest energy, knowing what that time is worth to me has really helped me to focus on how I spread myself across so many different things. It's also really helped me in deciding where it's worth spending money to effectively buy time. Let me explain.

There are endless things I could do with my time both professionally and personally. On the former that might be blogging, speaking, writing Pluralsight courses, doing workshops or simply meeting with people and building relationships. On the latter that's obviously spending time with family, watching a movie, going snowboarding and so on and so forth. Every activity draws down on my time and every one has its own reward. Some are monetary and immediate (running a workshop), others have longer term financial upside (Pluralsight courses), some build profile (blogging and speaking) and others reward in ways that are very hard to measure, such as playing with the kids. But each costs time and each pays something back, the trick is recognising this and prioritising appropriately. I've often sacrificed family time such as playing with the kids in order to work on the professional side because that provided higher value at the time. That may sound ruthless, but it enables me to spend more

time with them in other ways such as hitting the snow together for a week and focusing almost exclusively on them.

In terms of buying time, I happily pay housekeepers to visit, dry cleaners to take care of shirts (at least back in the day when I needed to wear them!) and people to wash the car. It's a low cost compared to what my time is worth, particularly if I'm paying for something that's tax deductable. A good example of this is that I'll pay for good seats on planes (or use my frequent flyer points) because it means I can do this:



I'm 6'5" and there is simply no way I can fit myself *and* my laptop in an economy sized seat. If I travel to Europe and back, there's 40 hours where I'll get zero work done (or very close to it) versus probably 20 hours of work and that's not including recovery time from not getting a good sleep either. That has a value and knowing what my time is worth helps me work out if there's an ROI in spending the extra money. I haven't always been able to justify that, it was only once there was sufficient reward on the effort invested that it made sense. There

are times where it still doesn't make sense, for example when ticket prices go astronomical or when I'm travelling with family and not working – I don't get the ROI on the ticket price then.

Tip: Figure out what return you're getting on your effort and use that to invest, prioritise and remove things that distract you from those priorities.

I optimise all the things

I used to spend a lot of time in racetracks trying to eke out every little bit of performance of the car. Brake a little later there, get back on the throttle earlier here and basically whatever it took to cut what usually amounted to no more than tenths of a second off a lap time. I'd incessantly analyse the repeatable things I did and look to optimise *everything* which in hindsight is a very agile-retrospective way of approaching things. It's also what I try to do in my daily life.

A while back I read a piece on why Zuckerberg always wears the same clothes. Here's the crux of it:

I really want to clear my life to make it so that I have to make as few decisions as possible about anything except how to best serve this community

Mark is optimising things to the extent that he's even trying to preserve the few brain cycles that would otherwise be devoted to picking his daily wardrobe in order to focus on *the things that actually matter*. Perhaps that's a bit eccentric, but I get it – I mean I get that the guy is making all these little tweaks and that they all add up to contribute to allowing him to focus more on what he does best.

A perfect example is using all available periods that are otherwise non-productive to do something useful. If I'm in a queue then I'm doing email. If I'm driving on my own then I'm listening to podcasts (the family aren't too keen on .NET Rocks!). If I'm walking back from taking the kids to school I'll try and

make phone calls.

I'm also very aware of when I'm becoming unproductive. For example, when a tweet or a notification distracts me. There's something that triggers and says "Hey, you're going off track, this is keeping you from what's important" and I try to adjust accordingly. Doesn't always happen of course, but at least I'm conscious of inefficiency.

Tip: Continue to perform self-retrospectives; how can you do this better? What worked? What should you do differently in the future?

I multipurpose absolutely everything I can

This one is a significant part of how I've done what I've done and I'm going to give you a heap of examples. When I wrote the <u>You're deploying it wrong series</u> on TeamCity, I was actually building out Pfizer's CI infrastructure. Writing the blog was the way I learned the ins and outs of TeamCity; it made me more effective in the office because I was publicising my views on the CI approach and opening them up to public scrutiny. I had to get things right in a way I didn't have to within the corporate environment. Partly that's because people were usually too polite to disagree (remember, it was the APAC region I looked after and culturally, you've very unlikely to be told if someone disagrees with you) and partly it's because I was the smartest guy in the room. Let me caveat that to try and avoid sounding conceited: pretty much everything at Pfizer was outsourced and the internal technical knowledge was gradually carved out (including with my departure) so bar one or two notable exceptions, there just weren't people there with the experience to voice an opinion. It's hard to debate the merits of build agents and MS Web Deploy with people who live in PowerPoint and Outlook! Incidentally, when I left and handed over management of the CI environment, it was that blog series that was my documentation – "Here you go

guys, here are the server names and everything else you need is on troyhunt.com". So I got public recognition for CI expertise, Pfizer got a great build environment, I made the handover a heap easier *and* I later got consulting work in the same space because of my public profile on it.

There are many, many other examples. A more recent one was the <u>Azure PowerShell blog post</u> I wrote earlier this year. Same deal as above in terms of the value of public scrutiny, but this time it was all part of trying to move 80 odd websites from a traditional hosting model into the Azure website and DB PaaS offerings. Not only did I write the blog post, but I also turned the whole thing into a Pluralsight course which mirrored exactly what I'd implemented in Pfizer – <u>Modernizing Your Websites with Azure Platform as a Service</u>.

In fact Pluralsight has been an excellent means of multi-purposing everything I do. Hack Yourself First was built based on many blog posts I'd written in the past and it remains one of my highest-paying courses. I've also done lots of conference talks by the same name and it's provided the framework for a very successful workshop I now do over and over again. The initial effort that I invested once to build the shape around that content has paid off time and time again in both a monetary way and in terms of expanding my profile and my influence.

A suggestion for anyone interested in following my approach: write a blog and write about what you're doing. That doesn't mean talking about the sensitive internal bits of your organisation, it means demonstrating knowledge about the stuff you'll have in your CV anyway. This is also relevant to my Ghost who codes blog post: by writing about what you're doing you'll do it more effectively, you'll help other people with the same challenges *and* you'll build your own profile. See how all that ties together so nicely? Oh – and don't let fear of corporate wrath stop you, there's a way to do this in a very mutually beneficial way and it would be a *very* rare case where there's a valid reason that can't be done. Check out <u>The Best Thing You Can Do for Your Career</u> on "side gigs" as well.

Tip: Work once, use many. I cannot over-emphasise this: Using the knowledge

you gain to create multiple things is massively important to productivity.

I plan the order in which things happen to maximise their effectiveness

A significant portion of what I produce happens in a well-planned sequence. I need to write something in order to have a baseline to refer to in something else or to form the basis of a new course or a talk I'm about to do or similar. I plan these things in order to maximise the value of each whether that be through more public exposure or improving my own knowledge.

I do also mix it up and that goes to that earlier point about task-switching between things that I feel like at the time. The post this week about <u>disabling password managers</u> was just a spur of the moment thing because I got an itch I needed to scratch. This post here has been in the works for months and it's going live now because it times in with my wife launching her blog yesterday (more on the significance of that later).

I also plan the timing of communications. I know the periods that are quiet on the web and I'd never post on the weekend or on an American holiday because both of those events take massive slices out of my audience. I re-share things at certain times of day because they maximise impact and all of these little things compound and help me extract more value out of the things I do.

Tip: Have some semblance of a plan in order to maximise the return you get on your effort.

I use a lot of hours each week to actually produce things

There are 168 hours in a week and I could never have done what I've done by using only a quarter of them for professional pursuits. It's not always a very palatable idea, but I had to use a big chunk of "non-work hours" to get to this point. I say that in quotes because the reality of my Pfizer job often meant early mornings and late nights courtesy of the time zones I worked across, so I had to make the time to pursue my aspirations outside of this.

The problem I often faced (and it's certainly not unique to me), is that most of my day job wasn't actually spent producing anything. I ran RescueTime for a while and found that I blew more than two hours a day in Outlook. That's just reading and writing emails too, it doesn't include all the time spent in other apps in order to respond to them. This wasn't actually producing anything, it was merely oiling the corporate wheels so that far enough down the chain somewhere somebody else could actually get something done. And it was enormously unfulfilling which is a large part of the reason I had to create my own things.

I often worked until 1am. I'd usually start at 6am. I *always* worked on weekends, albeit with leaving time for family activities as well. The laptop came on every holiday and I had to work very hard at balancing productivity with time out. Like I say, that's not very palatable to many people but this is what it took for me to get this amount of stuff done.

The bottom line though is simply this: I worked a hell of a lot of hours for many years and as much as all the previous points made a big difference to my productivity, I absolutely had to make dedicated time to work on these pursuits.

Tip: Time is an amplifier of the practices above. The more time you make to apply them, the more effective they are.

I get a lot of exercise and watch my

health carefully

This industry I'm in (and you probably are too) is not a healthy one. It's sedentary, it's bad for your posture and there's a good chance you'll cop a whack of RSI at some point. Particularly when there are high workloads like I've had, you've got to look after yourself. That means being both really cautious about what I eat and making sure I get plenty of exercise. I <u>talked to John Sonmez on his Get Up and Code podcast</u> last year and <u>wrote about some of the things I do to stay active</u> then.

I find certain types of exercise help me focus in different ways. When I windsurf, I'm often just out there for hours on my own and the mind wanders to big picture stuff. On the other hand, I regularly play high-intensity tennis by way of an hour and a half of non-stop drills in an evening and after coming home and getting cleaned up I can go to all hours of the night in ways I don't normally feel inclined to. Maybe it's endorphins helping things along, I don't know, but I do know that it contributes massively to my focus.

Even if it's only very mild, exercise is a great way of regaining focus. I'll regularly get up and go for a brisk walk to the shops just because it gets me away from the PC. Many problems have been solved just by changing context and clearing the head.

Tip: Take time out and do something as non-"sitting at the PC" as you possibly can – you need the mental break.

I have a supportive wife with a shared vision

I just can't emphasise this enough and I'll illustrate it by talking anti-patterns for a moment. So many times I hear people who want to get more involved in the sorts of things I've been doing say "I can't, my wife / husband / kids would go nuts". That's a perfectly fine position for the significant other to take and by no means is that a bad thing when that's your shared vision. It's when one party continually wants to head in a direction that the other doesn't support that things get tricky.

I had many years of working my arse off at all hours in order to reach this point and that put a lot of pressure on my wife. We have two small kids and she was working in a pretty high level job until recently too; my drive and ambition became her responsibility as well. And she supported me – *almost* unquestioningly – because ultimately, we share the same vision of how we want to live. It took many years for that effort to bear any fruit beyond my own personal sense of satisfaction which whilst important, doesn't do a lot for the family as a whole. It's only been over the last 18 months that she (and the kids, for that matter) have seen a return in a way in which we all benefit. Obviously an important part of that is financial, but it also translates into flexibility which means more time with the kids, more holidays and more generally just doing what we want, when we went. It's been her sacrifice as much as it has been mine.

And now I'm helping her to do the same, at least insofar as focusing on her public profile and building independence from corporate life in the way we've both previously known it. Just yesterday she's launched kyliehunt.com and will now start to be a lot more active on Twitter via kyliemhunt and other channels that will probably be familiar to many of you. If you appreciate what I've been able to do over the years, tweet her a quick thanks because boy does she deserve it!

Tip: Talk to your significant other about where you want to invest your time and how that benefits you collectively. Don't *not* do this and don't let it come between you; agree on where you're heading *together*.

Final thoughts

I've been as candid as possibly could in this post because many people were curious and I hope that it will actually help others to focus on doing wonderful things. This is not "the one true way" and much of it won't work for everyone; there may not be the same return on money spent, time may not be available and partners may not be as understanding. It's what's worked for me, make of that what you will and adapt accordingly.

I'll end on an observation that's really resonated with me:

85 percent of your financial success is due to skills in "human engineering," your personality and ability to communicate, negotiate, and lead.

Shockingly, only 15 percent is due to technical knowledge.

- Carnegie Institute of Technology

Whilst there's <u>debate about the origin and accuracy of this statement</u>, I've no doubt that my technical ability is but a small contributing factor to my success. Don't get me wrong – without it none of this would have happened – but the ability to communicate and influence others has contributed *significantly* to building the independence that I now have. Techie people are not renowned for these skills and if you can work on building those up, it'll put you at an enormous advantage.

That's been my story, I hope you found it valuable.

Comments

This is a great set of tips, but I have to say there was one jarring standout.

"I often worked until 1am. I'd usually start at 6am." This, with the surrounding paragraphs,

implies that you operated on an amount of sleep that sounds completely alien to me, even assuming you're sleeping through that whole off period. Most of the other things here are understandable tradeoffs (side projects vs recreation, family time vs work time, etc).

Following that standard would give me back 4 hours a day, a bigger productivity gain than anything else I can think of. But sleeping <5 hours with any regularity is off the table to me - even with quality sleep, it feels mind-eroding and health damaging.

Are you one of the \sim 1% of people who functions on four hours of sleep? Or is there some clever trick, or just a lot of tiredness?

Troy: I *often* - but certainly not always - had late nights and early starts. I try to make that the exception as I like getting a full 7 or 8 hours as much as anyone else. I tend to be driven by how I feel at the time; if I've been playing tennis until 9pm then I have the energy to work past midnight. Other times (like today), I wake up at 4am on my own (sometimes due to jet lag) so I get up and be productive.

IMHO, it's more about listening to your body and working within sustainable limits than it is forcing yourself to work tired.

This is a great post, thanks for sharing...it's the second time I've read it. You hustled so hard, for so long, before seeing huge returns. What motivated you to work so hard? Did you have a specific end-goal in mind that you were trying to achieve or was the destination more vague for you?

Troy: don't think there's any specific end-goal, at least not in the sense of "well that's all that done then"! I mean I've certainly had many personal objectives along the way and that's motived me to keep pushing, but I think a lot of it has been just enjoying seeing personal growth. Particularly when I compare this to corporate life where my growth was very much in their hands (seniority,

monetary, personal development), I'm relishing being in control of this myself now and that's a really major factor that keeps me going. Plus, I actually really enjoy what I'm doing:)

Thanks Troy. I keep re-reading this, there are just so many great tips for productivity.

But what if you don't already have an established level of success and exposure?

What are your tips for getting started down a path of "income security" and a great work-life balance? I'm guessing they all involve increasing one's public presence? Blogging, open source projects, making courses etc?

Troy: It's a journey Stuart, all of us started from zero profile and worked up from there. I re-read <u>my first ever blog post</u> recently and that's worth a read for context against this post here nearly six year later.

It's also worth checking out <u>The ghost who codes: how anonymity is killing your programming career</u> which I wrote when things were just starting to really take off for me. All of these things are interesting in terms of illustrating the journey and will hopefully give you some ideas around your questions.

Epilogue

Everything I wrote in that post is still on point today. In fact, just reading back over the headlines again now, every single one of them describes how I still live my life today. I love looking back at this post because it vindicates the views I formed at the time and I suspect I'll still feel the same when I look back at them in the years yet to come.

To look back on this post now is to look back on a previous life. An ex-life in an

ex-city with an ex-wife. The final point is the one that stings the most to reflect on, even after moving on to a new life in a new city with a new wife (to be). I could have excluded this post from the book and saved myself the pain of writing this epilogue, but then that would be ignoring an important part of my history and for what reason? To make things easier on myself? No, it's part of what made me who I am today and I'm very happy with that person, so this is an emotional scar I need to suck up and deal with. If anything, the points I made about relationships are more important to me now than ever, in particular the "shared vision":

"That's a perfectly fine position for the significant other to take and by no means is that a bad thing when that's your shared vision. It's when one party continually wants to head in a direction that the other doesn't support that things get tricky."

If I was to amend this today, I'd add that it's not just about sharing the same vision but maintaining it over time and being willing to sustain the level of commitment needed to reach it. That can be extraordinarily difficult, especially over the course of time as personal priorities change. Difficult, yet critical for a harmonious relationship.

HERE'S WHAT ASHLEY MADISON MEMBERS HAVE TOLD ME

This is possibly the most serious blog post I've ever written. For many people, Ashley Madison was nothing more than an adultery website inhabited by low-lives who got what they deserved after the massive breach of 2015. But for me, I saw a completely different side; one of pain and suffering, destroyed lives and in some cases, even suicides. I was seeing people at the absolute worst time of their lives and for some reason, they were telling me all about it.

For the most part, dealing with data breaches is just a mechanical process for me. Data comes in, I verify and load it, hit the go-live button then its job done and onto the next one. But not with Ashley Madison. I received an absolute outpouring of very personal messages from people in the breach, messages I was totally unequipped to handle. They showed a human side to data breaches beyond what I'd ever previously considered, and I really wanted the world to see that. Not to shame people, but rather to show that data breaches can have serious real-world consequences. I also wanted to reset the narrative from immediately assuming members of the site had some sort of moral deficit to showing that it actually drew in all sorts of different people.

I learned a lot from this incident, and I hope you do too in reading this post.

24 AUGUST 2015

found myself in somewhat of a unique position last week: I'd made the Ashley Madison data searchable for verified subscribers of <u>Have I been pwned?</u> (HIBP) and now – perhaps unsurprisingly in retrospect – I was being inundated with email. I mean *hundreds* of emails every day with people

asking questions about the data. Not just asking questions, but often giving me their life stories as well.

These stories shed a very interesting light on the incident, one that most people are not privy to and one that doesn't come across in the sensationalist news stories which have flooded every media outlet in recent days. When sent to me as an unknown third party in a (usually) foreign location, people tended to be especially candid and share stories that really illustrate the human impact of this incident. I thought I'd share some of those here – de-identified of course – to help people understand the real world impact of this incident and 'for those caught up in it to realise that they're among many others going through the same pain.

I responded to every legitimate email I received. Very early on I wrote up a Q&A and the following is the canned response I sent in response to almost every query:

My apologies for not being able to respond to you personally, I'm addressing questions of this nature via a Q&A you can find here: http://www.troyhunt.com/2015/08/ashley-madison-data-breach-q.html

Here's what Ashley Madison members have told me:



Lack of support from Avid Life Media

This probably shouldn't be surprising under the circumstances, but there wasn't much joy being had from concerned customers who wanted to get in touch with Avid Life about the incident:

I tried to reset the password and call them but they aren't answering phones or responding to emails

This is one of the things that struck me most about the entire incident – the *very* poor communication from Avid Life. At the time of writing, there has been no direct communication with members that I'm aware of, no notification on the front page of www.ashleymadison.com and in fact the site still talks about "discreet encounters", "trusted security" and "100% discreet service". The way they've handled this incident has been appalling – it's as if they've just stuck their fingers in their ears and sung "lalalalalala". And no, the legal action they've taken behind the scenes to track down the perpetrators and issue DMCA

takedown requestions does *nothing* to actually protect the impacted individuals. By now, we should have seen the usual offer of identity protection, admission of guilt and at least *something* to try and assist those who are having their lives torn apart by this. Instead there's nothing. Nada.

People aren't really concerned about their financial information

I found it odd that Avid Life Media felt compelled to issue a statement <u>that solely focussed on no financial data being compromised</u>. Do they *really* think that after the most intimate, private aspect of people's lives has been put on public display that a credit card their bank would simply replace if compromised is what they're worried about?! I had a *very* small number of requests like this:

How would I find out if any of my credit card info and/or email addresses have been breached? Thank you.

Even then, the requests about cards were thrown in with other queries about the data. Perhaps Avid Life made that statement to appease the <u>PCI</u> folks, but certainly card data is the last thing Ashley Madison members are worried about right now.

Lack of tech savvy

Those of us who live in technology often forget just how foreign it can be to those who don't. I've seen a lot of misunderstanding about fundamental technology concepts which victims of the breach obviously just haven't grasped:

My question to you is Ashley Madison has not responded to request for a

password change. So does that still get me the notification alert from you?

Now this website [redacted] if someone went to them and wanted to get my information & paid for this service. With having my email address. Could people get my information or would I get a notification from you stating that someone is requesting it?

I honestly found it hard to even understand some of these questions as the mechanics of databases and hackers and all sorts of other foreign concepts went over the heads of many people. That's totally understandable too and it just goes to show how everyday folks have been caught up in this mess.

Tor, BitTorrent and MySQL crash courses

Many people wanted to inspect the data for themselves, but with no knowledge of Tor or how torrents work (let alone the ability to then decipher the contents of MySQL scripts), most were left struggling:

Can I check any of this myself using a Tor browser, which I do not know how to use?

I have downloaded the data but I can't really make any sense of it, or in fact can't even open some of it up as its too large

I have downloaded the dumps, but I am not very handy so I'm not finding anything relevant at the moment. I own a Mac and I don't know how to open them, apart for using the standard txt editor and searching around.

I can totally understand the desire here but this simply isn't data that's consumable via your average person. Discovering it via Tor or downloading the torrent isn't particularly hard, but actually parsing the files and combing through the personal data spread across multiple tables is no simple task. For your average person, setting out to try and do this poses another risk altogether...

Falling victim to malware and other online scams

Following from the previous point, in desperation to find information, some people were resorting to downloading what they *thought* was the Ashley Madison breach, but evidently was something different altogether:

It seems easy to download the complete list from the pirate site. However the associated applications seem very dodgy

We always see this pattern: a serious international event happens (i.e. the recent Malaysia Airlines crashes) and immediately after we see nefarious individuals attempting to monetise either the pain of victims or the curiosity of onlookers. I've seen multiple sites purporting to offer the Ashley Madison breach which just require you to install this one little executable in order to view it...

Requests to search by fields other than email

I had a huge number of requests like this:

Is there a way u could search on my name if I gave to you

Will this data dump be eventually searchable by bill zip code?

I wanted to know if there's a way I can do a name search.

In some cases, people genuinely didn't know what email address they'd used. In other cases, I'll speculate and say that people were wanting to check up on other individuals which, of course, is precisely why I don't allow a search on HIBP by anything other than a verified email address. Searching by zip code is a perfect example – people don't want to do this to check their own exposure, they want

this feature to discover a range of people.

Data requests

One of the most frequent requests I got was to provide information on the actual data that had been exposed about the individual:

It's been so long I genuinely don't remember if I used a credit card, exchanged messages, what kind of personal information might have been in the profile, etc.

And i don't know if there is any point in asking you but can you tell me what information about me is in the dump?

I do not remember what was on my profile but am desperate to find out.

I am hoping to find out how much of my data is exposed and to prepare for the worst.

I just found out my husband's AM account is part of the hack. I want to know what information he put on the site.

Is there any way you can provide me with the info related to this email? At least then I can delete this email account and move on.

Now I am looking to confirm what I believe to be true so I can do damage control when the inevitable takes place. Some key info I want to find are:

- CC Txns (if any at all and corresponding date)
- Last Login
- Number of Logins
- Sign Up Date/Time*
- Cancellation Date/Time*

Is there no way you can tell me what info about me is on here? I've tried to

locate the data and cannot, I need to know how to prepare for this. Thanks

This is understandable – people want to assess their exposure – but I *always* declined not just because I simply couldn't do this for everyone, but because I have absolutely no desire to see personal information of this nature from Ashley Madison and then communicate directly with the impacted individuals about it.

Please erase me from the internet

You can understand the sentiment and for those who don't get how the web works, this would appear to be an entirely reasonable request:

I wonder if you could offer advice for trying to hide it again, take it off, remove it etc. or can this even be done?

Can i please unsubscribe my email so no one else can search me?

Do you know the reasoning why the company has not been successful in removing the material on Pastebin through the Digital Millenium law?

Could you assist in getting [redacted] off the AM dark web list?

As someone said to me in one of the comments on my blog, trying to remove your data from the web is "like trying to remove pee from a swimming pool". I added the DMCA comment in there as well because this has come up many times in the press. There's a good piece on it in an article that emerged after news of the attack first broke last month (paradoxically, stating that DMCA is the reason the full data hadn't been leaked), do read Parker Higgins' comment about the "fraudulent" use of the act in terms of its use for removing data breaches. Regardless, a US law will in no way stop the mass distribution of this data, particularly via a decentralised mechanism like torrents.

Can I please have the dump?

This was a common request:

Hi, can I get the bulk data dump for Ashley Madison can I trouble you for the tor page link?

It's an easy answer – no. At least you can't have it from me.

Payment records deanonymising members

Some people used non-traceable email addresses when signing up to the service, but then used their real identities in order to make payment:

My main fear is my credit card would be associated with the account at AM.

I used a burner email address but paid once for a full membership. Now it seems my name and address are affiliated with the breach.

My email is private, just for Ashley Madison. My real concern is, Is there any data which can be used to trace AM to me? For instance, I paid by a personal credit card when I first enrolled. How much trouble am I in?

Please please please delete that comment! It regards if cougarlife was hacked?! I dont know how to delete it... I think i accidentally logged into fb while posting when i thought u could be anonymous

That last one was from someone who commented on this blog using only a very common first name not linked to a profile but clearly the whole saga got them very worried about their own operational security. Obviously some members were conscious of protecting their identity in terms of hiding their membership, but didn't think through the digital footprints they leave by making online

payments. Whilst the payment files don't explicitly reference the identities in the membership database, both store the users' IP addresses, often allowing you to make implicit matches across the two.

The impact of public search services

Multiple services designed for *anyone* to search *anyone else's* email address quickly appeared and naturally, were quickly abused:

So got a call, from our church leaders yesterday, saying my husband's work email was on [redacted], oh my!

What. The. Fuck. I appreciate the curiosity that some people may have in terms of searching for other people they may know, but searching for groups of people within an organisation and for that organisation to be a *church* is unfathomable enough, but to then call up the spouse and notify them beggars belief.

Incomplete data on other search services

I was somewhat intrigued by messages like this:

Why does my email address--[redacted]--appear on yours but doesn't appear on three others, like [redacted] and [redacted]?

In fact I was so intrigued that I investigated it in more detail as the last thing I want is any inaccuracies in the HIBP data. What I found was that the two services mentioned in the above messages did not include some email addresses from the payment history files. This is alarming as it may be creating a false sense of security for impacted individuals and it just goes to show the responsibility

those of us standing up services like this take on board.

Closed email accounts and erasing the evidence

A lot of people were trying to effectively rewrite history by cancelling the email account they used for Ashley Madison. Either that or they'd legitimately moved on from both AM and the address they'd used for the site. Upon realising they needed access to the email account in order to search for it on HIBP, I got a lot of requests like this:

I used an alternate email address and have since canceled it out of sheer fear. How can i find out what, god help, if any of my info was leaked.

I had an email account [redacted] that I deleted in panic when the AM leak came out. I can see on other sites that it is included in the breach, but now that you've added the filter I can't see it on HIBP.

This account was closed when the business was closed down early last year as it went through a third company that supplied our web site at this time. Is there any way i can find out where the breach occurred ???

There was simply nothing I could do in these cases. Of course they could always search on another service which didn't require verification that they could access the email account, but certainly HIBP wasn't going to be able to help them out. The obvious problem here is that for all intents and purposes, "I don't have access to my old email account" is the same thing as "I don't have access to someone else's email account".

Accidental members

The observation has been made before, but the presence of a mere email address alone does not constitute infidelity on behalf of the account holder. When anyone can sign any email address up to the site, people who'd never even heard of Ashley Madison found themselves implicated:

I actually never signed up for this website which has lead me to believe that I have been victim of a scam. I have had numerous warnings of viruses on my computer. Perhaps this has something to do with it?

People like me are on the list despite NOT signing up on the website, because the website did NOT verify email addresses and someone gave mine as a supposedly fake address.

However, people seem to sign up for things all the time with my email address and I usually ignore it or do a quick password change on them so they have to move on.

Last night my wife asked me if I was one of the people that was using Ashley Madison. I haven't used the service but I know she's going to obsess about this so I did a search on a couple of sites where you could search email addresses for users. MY email address, this one, had a hit which is really perplexing to me since I've NOT used the service. Could someone have used my email address?

Of course these messages may also be ploys to convince their significant other that their presence on Ashley Madison was indeed none of their doing. The additional data attributes in the breach would tell the full story, which may also explain why I got so many data requests.

Suspicious wives

There's no question of gender equality here; *very* close to 100% of the emails I got were about men having accounts on the site. Understandably, there were

many suspicious wives asking me to check up on their husbands:

I wanted to know if you can search my husband's name/info for the Ashley Madison hack. I have found the AM site shown 2 times on his IPad history & a MILF hook up site when I looked at the history He claims they were "pop-ups" from porn sites.

That said, I have 20 years of my life invested with my husband & my gut tells me he is lying about it being on the Ipad & there are other things that lead me to believe he was a "member".

There's a lot of speculation about what the actual split between men and women on the site was (although I've not seen much on sexuality so am working on the assumption of predominantly heterosexual relationships), much of it relating to fake female accounts possibly created by Ashley Madison or accounts created by sex industry professionals to lure men into paying for services. It's all very conceivable and whilst we'll never know the actual numbers, I can say with great confidence that AM is *very* heavily male biased.

"Innocent" members

There is an assumption that those who signed up were always married and looking to have an affair. Whilst this is undoubtedly the case for many people, there was nothing prohibiting single individuals from joining the site:

HELP! I signed up for AM one night bit I'm actually single. I used my real email but fake info the rest of the process.

Whilst Ashley Madison may not represent the same moral high ground as other dating websites, there is a world of difference between someone in a committed relationship seeking out an affair and a single individual looking for a partner.

Alternate purposes for membership

Further to the previous point, there are other scenarios in which someone might create an account as well:

As a divorce attorney who often searched AM for my clients (and found a couple of cheaters there), I think it should be addressed that there are most likely women who merely joined AM as guests without paying or ever actually engaging- for the sole purpose of attempting to catch a cheating spouse.

I joined this site for 2 days about a year and half ago after my husband had an affair. I was having significant trust issues and joined ONLY to see if he was on the site.

You can't help but feel doubly sorry for these women; not only were they dealing with their husband having an affair, now they're also implicated as members of Ashley Madison themselves. It's a terrible situation to find themselves in and again, a poignant reminder that an email address on the site *does not* mean the individual intended to cheat on their partner.

Incorrect conclusions

An outcome I hadn't foreseen was some people thinking that *any* result for an email address on HIBP meant a presence on Ashley Madison:

Look dude, my wife want a divorce now since my email shows 'owned' when she put it in. Can you explain to her it's not for the Ashley Madison hack its checking the all pwned sites

This was actually for Adobe, the same breach I had three different accounts in!

Membership was from a different phase of life

We all go through phases of life where our views on things change. Many people have moved on from whatever that previous phase was, but now the Ashley Madison data is publicly haunting them:

Was a guest briefly some time ago. Different circumstances. Wanted to check now as life has changed and be sure.

I don't recall ever even visiting the site, but it's possible in some moment of general curiosity to see if people actually did that sort of thing.

Several years ago, when I was single (and recovering from a very bad breakup), I took out a profile on Ashley Madison

Not really worried as these are all old accounts from my single days but just curious as to what's floating around on the web.

I am single and not married, so this leak would make small harm, but it's a scary reminder of the perils of this new world we live in.

I was an AM member back when I was single and although technically shouldn't be concerned, my partner now is not one to take my word for it and will force me to sign up for notifications/verify my email and check my email.

I've included a few of these examples because I want to illustrate how important it is not to immediately assume that everyone on the site is cheating on their partner even if they were legitimate, paid up members. Of course many are (or at least "were"), but it's important not to immediately make assumptions just because someone's email address was on the site. Others will pass their own moral judgement on whether individuals should be registering on a site primarily designed for sexual encounters, but let us not confuse that with the

issue of adultery where another innocent party is adversely affected.

It was never really serious...

I found people frequently justifying their account to me, as if they worried that a stranger on the other side of the world might judge them:

I know you're not judgmental, but I'd be remissed if I didn't state that I never actually met anyone - it was more of a game to see how i could get responses.

Never did anything but look around and deleted in like 2010. Really sad and scary.

Long story but was not cheating at all but had a profile created and then paid to have it deleted with their pay to delete function.

I joined Ashley Madison one night bored, honestly. Used my real email, but fake info from there on and never used a CC or got a real membership.

Spent 15 mins and have never been back

I've been caught up in it, my own story a drunken evening, curious about the site, signed up, thought, OMG this is not a good thing to do, got out of the site, never touched it again

If we take these messages at face value – and I'm not sure there's really much value in lying privately to a stranger for no apparent upside – many people were indeed just curious. Of course some people could be fabricating the message, but it's entirely feasible that no nefarious activity actually took place.

Remorse

It shouldn't come as a surprise, but there was a huge amount of this:

No question I made a terrible, terrible mistake and pray to god this doesnt come out and ruin my family.

I am not married but Ashley Madison was/is a mistake I made and wonder how much risk I am at being publically embarrassed and more importantly embarrassing my Parents and Siblings.

I feel pretty sick and foolish - I've done nothing other than a few two sentence chats but I still don't want to have to deal with this.

Last night was the worst night of my life. Found out my AM account had been breached.

I regret having signed up to the site and now terrified about hurting those around me, especially the one I love.

I am absolutely sick. I can't sleep or eat and on top of that I am trying to hide that something is wrong from my wife.

My wife found out about it after I had exited the site and we have gone through a long period of working on our relationship. Its been a long and painful journey - but a private one - and we are closer than ever before, and I bitterly regret what I did.

These were often very raw emotions and as the comment above says, it's a private journey for many people. Regardless of your take on the ethics of someone being on the site in the first place, most people would agree that in situations like this, the individuals deserve the privacy to work on their relationships and move forward in life. This incident will seriously jeopardise the ability for many couples to do just that and unfortunately the prevalence of publicly searchable AM databases merely fuels that fire and sets these couples back even further.

Fear and desperation

Clearly many people were fearful of being discovered for having an account on the site, either by their partner or by other members of the community. The fear of potential consequences often came through in a very raw way:

I love her very much and don't want to lose her, I am deeply worried that she will leave and greatly impact my life.

I literally cannot sleep and never met anyone but am terrified as what might happen.

I never met anyone on the site, I'm not married, but this has me spinning. I need advice. Please help.

At this point I'm desperate. Worried that something like this could ruin my life/marriage when I was not on that site for anything that I can remember, possibly curiosity/joking with friends, but I can't recall. I've barely slept over the past day due to worry

This while situation is very confusing and scary.

My stress levels are through the roof, still hoping that by some miracle this will just be forgotten about and no one will want to search me up.

My last resort is asking you if you could PLEASE PLEASE PLEASE help me out and let me know what you have on me.

Sorry, I appreciate that must sound like a completely naive/desperate question, but that's the level I'm playing at.

What would be impossible to explain away - and what I would most feel guilty about - is the very detailed personal intimate information about my wife shared with strangers during my 'erotic' chats.

Admittedly, it was hard to read comments like the last one and not feel

resentment. Having that canned response available and merely directing people to the Q&A saved me from having to construct very difficult personal responses to emails like this. But do take the other ones on board too; this is the real world consequence of this event.

The impact on families

As a father myself, the hardest messages to read were the ones like this:

But I'm just a guy here with a wife that I really do love, I regret what I did, and I have two beautiful kids that will get sucked int to this too. Its just horrible.

I have couple of 3 year old kids. I can tell you my amount of activity on these site was basically limited to one or two session logins and more of just curiosity on what's there.....And in this case, looks like curiosity could kill the cat.

Tell your wife and kids you love them tonight. I shall do the same as I really don't know if I will have many more chances to do so.

I read that last one right before going to bed last night and it was difficult to grasp; extramarital affairs tear families apart. You don't need Ashley Madison for that to happen and arguably the guys making these comments deserve to go through some degree of pain, but you can't escape the human tragedy that this data breach has brought to a head. It's hugely distressing not just for the members who did indeed have affairs, but their families as well.

Real world consequences

It's not always obvious just what impact a presence on Ashley Madison can have

in "the real world", I certainly learned things I was never expecting:

adultery is a punishable offense under the U.S. Army's Uniform Code of Military Justice, and while simply having an active account at this website doesn't indicate any wrongdoing, it's possible that as the data become more publicized, some people are in for a lot of headaches.

One of the big concerns has always been that someone will take their life as a result. Allegedly, this may have already happened and it's hard to see how it wouldn't happen with such a huge user based impacted by such a significant event on so many lives.

Impact on professional life

A number of people were really worried about what membership of Ashley Madison – regardless of their context in there – might mean for their professional career:

How can this show up in a back ground check for jobs or anything if I have and provide this new email account to the admission boards and employers?

How do I keep it private from clients, customers, relatives etc.

I would like to know as I am very concerned but the whole mess and am a school teacher and really want to know what information they will eventually have access to.

And now my email address (which is my actual email address...dumb) is available to anyone who searches it. I am a professional and this could potentially be devastating.

In an era where employers are increasingly focused on building profiles of potential hires, I totally understand the concern. There's a good example of this concern in the public comment thread of my first Ashley Madison post and you

can sense the trauma this is causing the woman. That thread also demonstrates that whilst this is never something that *should* be used against someone seeking employment, the reality is that it will become one more data attribute in the increasingly rich profiles that are built up about individuals. There will surely be those that pass judgement against members *regardless of their context on the site*, let me give you some examples.

They got what they deserved

I want to add this here after all the other comments to illustrate how shortsighted some people are being about the breach. If you've read through all the comments above you would have seen many different levels of involvement in the site from entire innocence through to outright betrayal. Yet somehow, there are those who seek to tar everyone with the same brush:

JUSTICE for all the good people getting cheating on. Im glad the list has been exposed.. I don't care if other innocent people that weren't cheating were exposed that's the risks you get when signing up for this crap online TOO BAD.

If you ended up using an email address that you've shared with anyone else, you deserve to have your information exploited in such a way.

the fact that 30 million sleazebags had their identifies and details revealed by these hackers fills me with amusement more than horror. The only improved result to my mind would have been a letter addressed to their home addresses with ASHLEY MADISON membership update printed in large letters on the front.

The chickens come home to roost. I'm glad someone is providing some true justice in the world. It sucks to be cheated on and I hope everyone on that site feels like shit and loses someone who truly cared for them.

Anyone who signed up to this sick site deserves everything they have coming to them.

These are largely from public comments made on posts such as <u>my original one</u> <u>on how I'd handle the data breach</u>. I hope this offers some perspective to those who wish to pass blanket moral judgements on everyone. As much as Ashley Madison's mission statement is centred around the premise of infidelity, this incident is far more complex than just a bunch of cheating spouses.

In summary

This has been a lengthy post as I've continued to add to it as the messages have flooded in. I've been very careful to choose only messages that disclose nothing of the sender and this has meant not sharing the vast majority that came in. If nothing else, I hope it demonstrates how much of an impact this is having on lives, both those who set out to cheat on their spouses and the innocent bystanders be they accidental members, curious onlookers or the partners of those who have been outed. This incident needs to be approached with the understanding that for many people, this is the worst time of their life and for some, it feels like the end of it.

Comments

I work for IT Security for a large multinational corporation. Our concern was that if anybody had signed up using a corporate email address, then they could be targets for blackmail. After acquiring a list of email addresses sign up for the site, I searched for signups using our domain.. there were quite a few. Some of those addresses were clearly invalid, and one was an info@ address.

Those addresses were checked to see which ones were currently valid. The job roles of those

individuals was checked, and the list of roles ONLY was passed up the reporting chain to evaluate the risk which was deemed minimal. Only one person in the corporation has the list of names, the other parties involved in risk management just have the generalised job titles. The parties with email addresses listed have not been informed. In addition, there has been some discrete email monitoring to ensure that no email threats have been received.

In the view of the corporation, we do not prohibit the use of "dating" sites in the employee's own personal time. We discourage the use of corporate email addresses for personal use, although it is possible that these sign ups could be a decade old or more, and of course there is no guarantee that they are genuine.

Had we discovered a person we thought to be at risk of blackmail, then the next steps would be to go through our legal department for guidance.

Troy: That's a very insightful, well thought out and responsible response, thank you so much for sharing it. I've heard from a number of people that use HIBP's domain notification service and they were in a quandary as to what to do once they knew who was impacted. I'll point them to this comment - well done!

Epilogue

For years after the Ashley Madison breach, I still had people forwarding on extortion emails they'd received. *Years!* The tail of this breach is a long one and to this day, I can't think of an incident that has had a greater real-world impact on people.

This post also became one of my most commented ever with hundreds of people sharing their views. Many were looking for answers, others were using my blog as a channel to pour out their emotions. Reading back through some of those comments now, seeing people say things along the lines of "just wanting to end it all", is horrifying. But equally, I was fascinated that people

were choosing to confide such serious thoughts in me and my audience. Anonymously, of course, but then they thought they were anonymous on Ashley Madison too, so I'm still surprised to see such personal thoughts shared on troyhunt.com.

It also didn't prove to be the "watershed moment" I saw many people predicting. We didn't "fix" data breaches after this. We didn't get better at preventing them. We didn't collect less new data or purge more old data, everything just escalated. Even as privacy laws such as GDPR and CCPA came into effect, the rate of data collection and consequently data breaches just accelerated. I mention that here because this was one of those incidents where many people questioned if it was finally the one that was bad enough to cause the industry to do better. Turns out it wasn't.

For a while there, it looked like they were trying to become a classier adultery site. Take a quick spin through archive.org and you can see that in mid-2015, they were all "Life is short. Have an affair." with the image of a sexy woman, finger in front of mouth whispering to the audience whilst wearing her wedding band. Fast forward a year and now it's "Find your moment" with an artsy image of a much more, uh, "well dressed" woman looking coyly at the camera. Today? "Life is short. Have an affair." and the original girl is back too. That slogan is a registered trademark too, as is "When monogamy becomes monotony" which appears a bit further down the page. Apparently, sex still sells.

In 2021, I started getting contacted about Ashley Madison for a different reason - a documentary is going to be made on it. Multiple different production companies reached out, obviously pitching for the gig to produce the documentary. Eventually, Minnow Films appears to have been successful, something I believe I can share here as myself and others involved in the incident were asked to do some Twitter outreach in April 2022 to see if any victims would like to come forward. I'm really curious to see who actually wants to be featured on a "premium, global streaming platform", as they referred to the (unnamed but I suspect kinda obvious) company that will broadcast the

show. It seems there's life left in this story yet.

CONTROLLING VEHICLE FEATURES OF NISSAN LEAFS ACROSS THE GLOBE VIA VULNERABLE APIS

There's always excitement when finding a vulnerability like this. As you'll read in the blog post, it wasn't actually me that found it, but a student in one of my workshops that took the techniques we discussed, went back to his hotel room then enthusiastically arrived at the class first thing the next morning with his finding. It was exciting sitting there with him poking away at the vulnerable API, just waiting to see what would happen next. Keep that in mind as you read through this; a person with **no prior experience** in this sort of thing whatsoever was able to find a major, glaring vulnerability in software controlling features of a car that had been rolled out to probably tens of thousands of people. That'll give you a sense of how low the bar is set, not just for this incident, but so many others I've dealt with since.

24 FEBRUARY 2016

ast month I was over in Norway doing training for <u>ProgramUtvikling</u>, the good folks who run the NDC conferences I've become so attached to. I was running <u>my usual "Hack Yourself First" workshop</u> which is targeted at software developers who'd like to get up to speed on the things they should be doing to protect their apps against today's online threats. Across the two days of training, I cover 16 separate discrete modules ranging from SQL injection to password cracking to enumeration risks, basically all the highest priority security bits modern developers need to be thinking about. I also cover

how to inspect, intercept and control API requests between rich client apps such as those you find on a modern smart phone and the services running on the back end server. And that's where things got interesting.

One of the guys was a bit inspired by what we'd done and just happened to own one of these – the world's best-selling electric car, a Nissan LEAF:

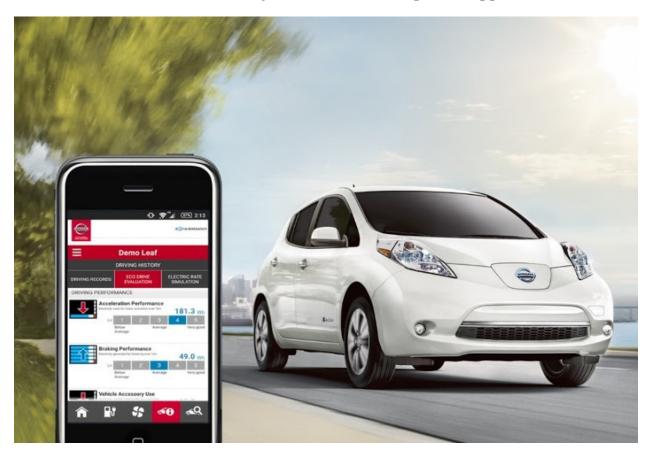


What the workshop attendee ultimately discovered was that not only could he connect to his LEAF over the internet and control features independently of how Nissan had designed the app, he could control other people's LEAFs. I subsequently discovered that friend and fellow security researcher Scott Helme also has a LEAF so we recorded the following video to demonstrate the problem. I'm putting this up front here to clearly put into context what this risk enables someone to do then I'll delve into the details over the remainder of the post:

We elected for me to sit outside in a sunny environment whilst Scott was shivering in the cold to demonstrate just how remote you can be and still control features of someone else's car, literally from the other end of the earth. Following is a complete walkthrough of the discovery process, how vehicles in other countries can also be controlled and a full disclosure timeline of my discussions with Nissan.

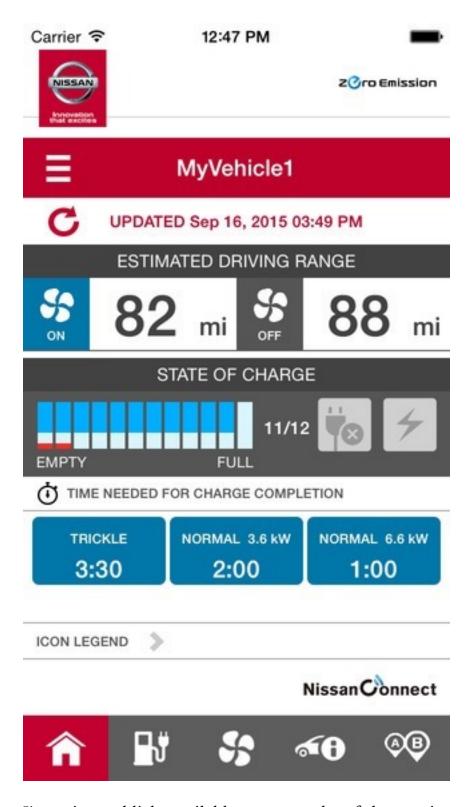
Connected LEAFs

The LEAF is an electric car which is particularly popular in countries like Norway which offer massive financial incentives to stay away from combustion engines. It does all the things you'd expect of a modern EV and because it's here in the era of the internet of things, it also has a companion app:



Back at my workshop in Oslo and being the curious type, Jan (not his real name – he requested to remain anonymous) goes back to his hotel after the first day of

the course and <u>proxies his iPhone through Fiddler running on his PC</u> as we'd done during the day (this was on January 20). This takes a few minutes to setup and effectively what it means is that he can now observe how the mobile app talks to the online services. Jan then fires up the <u>NissanConnect EV app</u>:



I'm using publicly available screen grabs of the app in part so as not to disclose personal information about Jan and in part because his app runs in Norwegian. When the app opens, he observes a request like this (I'll obfuscate host names and the last five digits of VINs throughout this post):

```
GET https://[redacted].com/orchestration_1111/gdc/BatteryStatusRecordsRequest.php?RegionCode=NE&lg=no-NO&DCMID=&VIN=SJNFAAZE0U60XXXXX&tz=Europe/Paris&TimeFrom=2014-09-27T09:15:21
```

Which returns the following JSON response:

```
{
    status: 200,
    message: "success",
  - BatteryStatusRecords: {
        OperationResult: "START",
        OperationDateAndTime: "jan 21,2016 21:47",
      - BatteryStatus: {
            BatteryChargingStatus: "NORMAL CHARGING",
            BatteryCapacity: "12",
            BatteryRemainingAmount: "12",
            BatteryRemainingAmountWH:
            BatteryRemainingAmountkWH:
        },
        PluginState: "CONNECTED",
        CruisingRangeAcOn: "135664.0",
        CruisingRangeAcOff: "157904.0",
        NotificationDateAndTime: "2016/01/21 20:47",
        TargetDate: "2016/01/21 20:47"
    }
}
```

This is pretty self-explanatory if you read through the response; we're seeing the battery status of his LEAF. But what got Jan's attention is not that he could get the vehicle's present status, but rather that the request his phone had issued

didn't appear to contain any identity data about his authenticated session. In other words, he was accessing the API anonymously. It's a GET request so there was nothing passed in the body nor was there anything like a bearer token in the request header. In fact, the only thing identifying his vehicle was the VIN which I've partially obfuscated in the URL above.

The VIN is the <u>Vehicle Identification Number</u> which uniquely identifies the chassis of his LEAF. It is by no means a "secret" suitable for authorisation purposes, the significance of which I'll come back to shortly.

On the surface of it, it looked like anyone could get the battery status of Jan's vehicle if they knew his VIN. Not ideal, but not exactly serious either as it's a passive query (it doesn't actually change anything on the vehicle) and there's also nothing of a personal or sensitive nature returned in the response beyond potentially telling you when it was last driven based on the OperationDateAndTime field. So Jan kept looking.

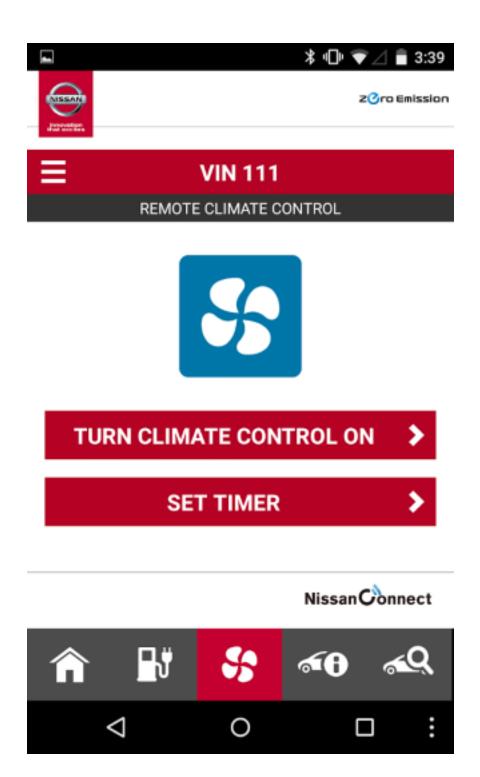
He found he could check the status of the climate control using this request:

```
GET https://[redacted].com/orchestration_1111/gdc/RemoteACRecordsRequest.php?RegionCode=NE&lg=no-NO&DCMID=&VIN=SJNFAAZE0U60XXXXX
```

Which then returned a similar status result:

```
{
    status: 200,
    message: "success",
  - RemoteACRecords: {
        OperationResult: "FINISH",
        OperationDateAndTime: "jan 22,2016 08:39",
        RemoteACOperation: "START",
        ACStartStopDateAndTime: "jan 22,2016 08:39",
        CruisingRangeAcOn: "134400.0",
        CruisingRangeAcOff: "159040.0",
        ACStartStopURL: "",
        PluginState: "CONNECTED",
        ACDurationBatterySec: "900",
        ACDurationPluggedSec: "7200"
    },
    OperationDateAndTime:
}
```

This is reflected within the app on this screen:



But again, it's passive data – is the climate control on or off and as a result, what should the buttons say. But then he tried turning it on and observed this request:

GET https://[redacted].com/orchestration_1111/gdc/ACRemoteRequest.php?

That request returned this response:

```
{
    status: 200,
    message: "success",
    userId: " ",
    vin: "SJNFAAZE0U60 ",
    resultKey: " "
```

This time, personal information about Jan was returned, namely his user ID which was a variation of his actual name. The VIN passed in the request also came back in the response and a result key was returned.

He then turned the climate control off and watched as the app issued this request:

```
GET https://[redacted].com/orchestration_1111/gdc/ACRemoteOffRequest.php?RegionCode=NE&lg=no-NO&DCMID=&VIN=SJNFAAZE0U60XXXXX&tz=Europe/Paris
```

All of these requests were made without an auth token of any kind; they were issued anonymously. Jan checked them by loading them up in Chrome as well and sure enough, the response was returned just fine. By now, it was pretty clear the API had absolutely zero access controls but the potential for invoking it under the identity of other vehicles wasn't yet clear.

Connecting to other vehicles

When Jan came into the workshop the following day, he also brought in a picture he'd managed to locate by searching the web:



This was the vehicle's VIN which clearly, left us curious (obfuscation is mine, it's legible in its entirety on the web).

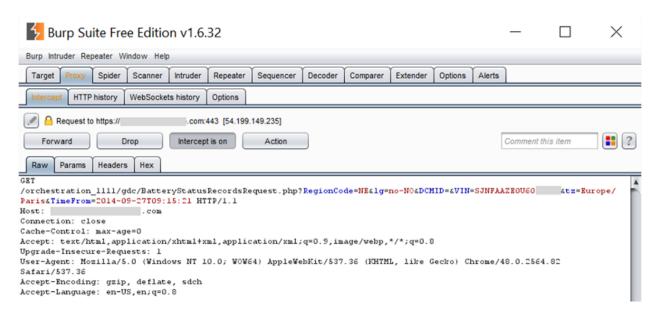
Let me clarify something before going any further and it's something I harp on about in my workshops too; when a potential security flaw is identified, you've got to think very carefully about how you proceed with verification. You need to have a sufficient degree of confidence that it's a legitimate flaw before reporting it ethically (which is what we ultimately did), but you also need to ensure you don't breach someone else's privacy or impact them adversely in any way. We wouldn't, for example, want to start operating mechanical features of someone else's car such as turning on the climate control nor would we want to retrieve personal information about them, even if it was just their username.

The VIN above differed merely by the last 5 digits. We grabbed the number and plugged it into the request to get the battery status – a request that didn't change anything nor disclose anything private – and got this response:

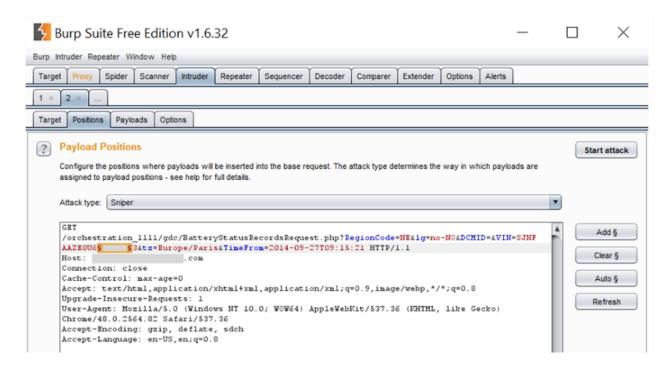
```
status: "-5035",
  message: "Not Specification GDC [TCUID:]",
  ErrorCode: "-5035",
  ErrorMessage: "Not Specification GDC [TCUID:]"
}
```

This appeared to indicate that the response couldn't be processed but it wasn't clear why. On reflection, it's possible that the VIN hadn't been registered for the app. It could also be possible that one of the query string parameters in the first URL I shared above wasn't valid for that VIN. For example, the RegionCode field may not have matched with the vehicle's location. Without a positive result from the API, we couldn't emphatically conclude that there was indeed a lack of authorisation.

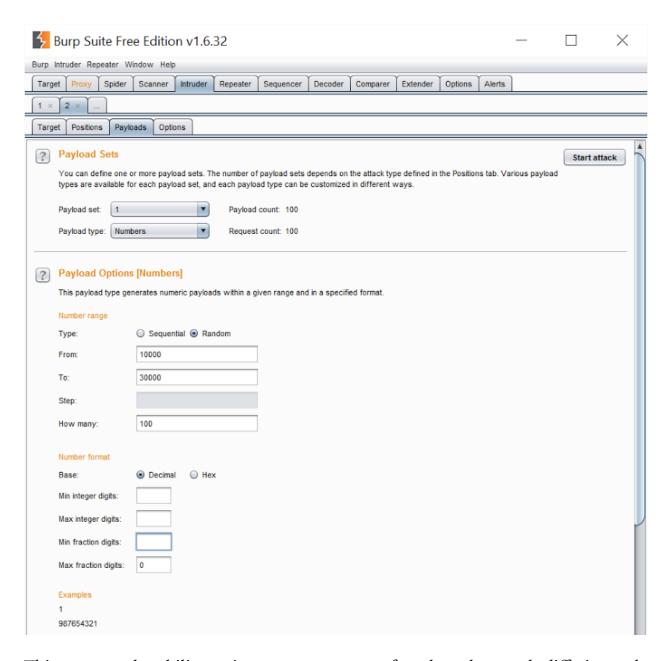
The thing about VINs though is that they're easily enumerable. Both Jan's and the VIN found on the web were identical except for the last 5 digits which meant we could easily test for other matches using a tool like <u>Burp suite</u>. We proxied Chrome through Burp then issued the battery status request again:



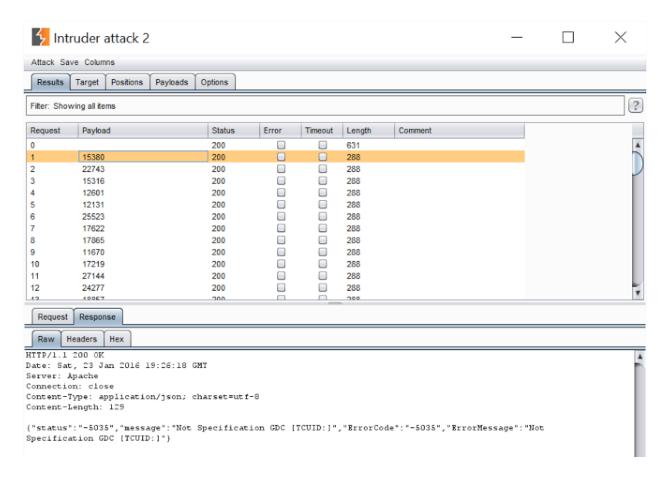
We then sent it over to the Intruder feature and added one position for payload insertion:



This was the last five digits of the VIN, those being the ones which differed across both Jan's and the number found online. (Note: not all LEAF VINs necessarily differ by just the last 5 digits, the VIN specification allows for the range to be broader, i.e. it may be the last 6 digits. Our test simply kept the range constrained between known numbers for the sake of time.) We then configured Burp to randomise those last 5 digits and choose integers between 10,000 and 30,000 which is the range both Jan's and the VIN online fell within:



This gave us the ability to issue requests one after the other, each differing only by a unique VIN in the payload column. We didn't need to test all 20,000 possible VINs within that range, we just had to issue requests until we found one that returned the battery status of another vehicle. We started Burp issuing the requests:



Request 0 in the screen above is the one to Jan's car which returned a response size of 631 bytes. The subsequent responses with the randomised VINs mostly returned 288 bytes and the response you see in the screen above. Until we found one that didn't:

```
{
    status: 200,
   message: "success",
 - BatteryStatusRecords: {
       OperationResult: "START",
       OperationDateAndTime: "jan 18,2016 22:05",
     - BatteryStatus: {
            BatteryChargingStatus: "NOT_CHARGING",
           BatteryCapacity: "12",
            BatteryRemainingAmount: "8",
           BatteryRemainingAmountWH: "",
            BatteryRemainingAmountkWH: ""
        },
       PluginState: "NOT_CONNECTED",
       CruisingRangeAcOn: "87792.0",
       CruisingRangeAcOff: "99696.0",
     - TimeRequiredToFull: {
            HourRequiredToFull: "8",
            MinutesRequiredToFull: "30"
       },
     - TimeRequiredToFull200: {
           HourRequiredToFull: "5",
           MinutesRequiredToFull: "30"
        },
     - TimeRequiredToFull200 6kW: {
            HourRequiredToFull: "3",
           MinutesRequiredToFull: "0"
        },
       NotificationDateAndTime: "2016/01/18 21:05",
       TargetDate: "2016/01/18 21:05"
   }
}
```

This wasn't Jan's car; it was someone else's LEAF. Our suspicion that the VIN was the only identifier required was confirmed and it became clear that there was a complete lack of auth on the service.

Of course it's not just an issue related to retrieving vehicle status, remember the other APIs that can turn the climate control on or off. Anyone could potentially enumerate VINs and control the physical function of any vehicles that responded. That's was a very serious issue. I reported it to Nissan the day after we discovered this (I wanted Jan to provide me with more information first), yet as of today – 32 days later – the issue remains unresolved. You can read the disclosure timeline further down but certainly there were many messages and a phone call over a period of more than four weeks and it's only now that I'm disclosing publicly, right after I received an email from a Canadian follower...

Vulnerable LEAFs in Canada

By pure coincidence, just as we hit the four-week mark since initial disclosure and I was about to revert to Nissan yet again, an email landed in my inbox from a Canadian follower titled "weird Nissan api". It started out like this:

I read your Vtech article and though that you would be well placed to appreciate this.

Im a Nissan Leaf owner and I found out that Nissan security is pretty abismal. They have an App to remote start charging, start/stop the AC/ Heat, and get updated on current state of the vehicule. http://itunes.apple.com/ca/app/nissan-canada-leaf/id450031231?mt=8

This came in just last weekend on 20 Feb and it went on to explain the following:

I found out that the whole API is unauthenticated and only require the VIN to target a vehicle. To add insult to injury those action are from simple http

Get request.

details on how to: (site in french) http://menu-principal-forums-aveq.1097349.n5.nabble.com/Nissan-Canada-Leaf-Carwings-td37239i20.html#a38494

This is precisely what Jan had found in Oslo and what's more, it was being discussed openly on a forum. Browsing through the discussion courtesy of Google translate, clearly people were not happy with the Nissan app. In fact, they were so unhappy that one post suggests taking the app out of the picture altogether and controlling the vehicle's functions by making requests directly to the APIs:

For hard-core, the following information:

URL to activate / deactivate air conditioning / heating. Put your VIN in
the URL, this works very well in your browser. Create bookmarks with
these 2

https://canada.nissanconnect.com/owners/leaf/setHvac?
vin=1N4A.....5520&fan=on
https://canada.nissanconnect.com/owners/leaf/setHvac?
vin=1N4A.....5520&fan=off

They go on to conclude precisely what we had earlier on:

In all this, it works for me without being authenticated, which is very surprising, and not safe at all, this means that anyone can act on any vehicle, provided it knows the VIN (in more is it not written down the visible windshield everyone?). Looks like the authentication uses has get the VIN in the user profile.

Now this was back in December so we're talking a couple of months ago already. Note also that the URLs above are different to the API endpoints we saw for the Norwegian instance (I obfuscated the other host names as I've not seem them discussed publicly, but even the paths are different). It's an odd design decision

for a global car manufacturer to segment their app in this way. There are always local idiosyncrasies to be considered (particularly in the auto industry), but there appears to be very little reuse across Canada and Norway in terms of how the API is implemented. It has me wondering if perhaps the build of these apps is delegated to local groups who perhaps don't pass through the same levels of rigour you'd expect at the global level.

The person who reported the Canadian finding to me finished up by saying this:

My hypothesis on this is that it was bound to surface due to the poor quality of the app, the more tech savvy "with free time" users will thinker with broken things to get them working for them. The fail was probably discovered soon after the app change and multiple times but by people that didn't fully appreciate the greater implication or by people like me that didn't know what to do with that knowledge.

His first sentence is spot on – the ease of discovery of this risk is high as is evidenced by three separate parties already finding it independently (my Norwegian student, the Canadian follower and the folks in the forum). The Norwegian case alone was cause for concern and the Canadian one showed that the issue was now well and truly out there in the public domain, but I wanted further verification which is where Scott Helme came into the picture.

Nissan LEAFs in the UK

It was by pure coincidence that <u>Scott Helme</u> from the video in the intro has a LEAF, that he's a security professional and that I spent some time with him in the UK just after my Oslo workshop. It was then that the penny dropped and we both realised that he could be of assistance. He's proofed everything I've written here and obviously also offered up his own car to verify that indeed, it's only the VIN required to operate the functions described in this blog post. Given his involvement, I asked him if he wouldn't mind sharing his own view of the

situation and he gave me this paragraph:

Fortunately, the Nissan Leaf doesn't have features like remote unlock or remote start, like some vehicles from other manufacturers do, because that would be a disaster with what's been uncovered. Still, a malicious actor could cause a great deal of problems for owners of the Nissan Leaf. Being able to remotely turn on the AC for a car might not seem like a problem, but this could put a significant drain on the battery over a period of time as the attacker can keep activating it. It's much like being able to start the engine in a petrol car to run the AC, it's going to start consuming the fuel you have in the tank. If your car is parked on the drive overnight or at work for 10 hours and left running, you could have very little fuel left when you get back to it... You'd be stranded

Of course the other thing we covered in the video was pulling the driving history from the vehicle which looks like this:

```
TargetDate: "2016-02-21",
PriceSimulatorDetailInfoTripList: {

    PriceSimulatorDetailInfoTrip: [

       - {
             TripId: "1",
             PowerConsumptTotal: "78.76",
             PowerConsumptMoter: "87.42",
             PowerConsumptMinus: "8.66",
             TravelDistance: "146",
             ElectricMileage: "1.2",
             CO2Reduction: "0",
             MapDisplayFlg: "NONACTIVE",
             GpsDatetime: "2016-02-21T11:26:49"
         },
             TripId: "2",
             PowerConsumptTotal: "332.47",
             PowerConsumptMoter: "501.99",
             PowerConsumptMinus: "169.52",
             TravelDistance: "1542",
             ElectricMileage: "2.9",
             CO2Reduction: "0",
             MapDisplayFlg: "NONACTIVE",
             GpsDatetime: "2016-02-21T11:31:04"
         },
             TripId: "3",
```

These are two trips he took on Feb 21 when he dropped his son off at his parents. He took two other trips that day, one to go snowboarding and another one to return. They were all recorded by the vehicle and are publicly accessible if you know his VIN which again, is displayed in his windscreen or can simply be

guessed by enumerating through those last five digits.

This gives you details on movements per day which raises all sorts of privacy risks. Scott gave me a comment on that too.

The other main concern here is that the telematics system in the car is leaking *all* of my historic driving data. That's the details of every trip I've ever made in the car including when I made it, how far I drove and even how efficiently I drove. This could easily be used to build up a profile of my driving habits, considering it goes back almost 2 years, and predict when I will be away from home. This kind of data should be collected and secured with the utmost respect for my privacy.

Whilst it's not specifically personally identifiable information such as the individual's address, by the time you have a VIN which you know belongs to a LEAF registered within a specific country, it may not take too much effort to fill that gap. For example, down here in Australia we have services such as revscheck.com.au which can report on a pretty extensive set of data based on nothing more than a VIN. Jan sent me an equivalent service for Norway at vegvesen.no. I suspect that there are multiple other avenues where additional data about the vehicle and the owner can be retrieved once the VIN is known and that opens the door to a raft of other possible privacy risks.

Rectifying the risk and opting out of the service

The underlying risk is simple and I'll quote Scott's comments on it:

This API thing is just nuts. It's not even like they just missed auth or didn't check, it's actually not implemented.

It was built, intentionally, without security...

Clearly the answer is to implement appropriate authorisation on all API calls, which when building an app in the first place would be a trivial feature to add. It's trickier to add to a "brownfield" app though and in Nissan's case, even trickier again because of the design of it. What's unique about their approach is that Norway and the UK seem to be hitting a completely different set of APIs to Canada. In fact the European API is on a host not even owned by Nissan, it's registered to "ZENRIN DataCom CO.,LTD." which may mean that we're looking at multiple API endpoints controlled by different parties that need to be rectified. Then of course the apps have to be updated across different client devices (iOS, Android, etc) and for different languages then pushed out to consumers. Whilst waiting for this to happen, LEAF owners remain at risk.

Because the question would inevitably arise, I asked Scott how he'd opt out of the service and he provided some steps:

Given the ease with which someone can enumerate valid VIN numbers, this issue raises a few concerns. It could be a huge inconvenience to have someone run my car flat by using the heating and accessories all day and the exposure of my entire driving history poses quite the privacy concern too. To disable CarWings, owners need to login to the service form their browser, it can't be done through the mobile app. Once logged in, select 'Configuration' from the menu and there is a 'Remove CarWings' button. It appears to be greyed out but the button does work. Once clicked you will receive a prompt to confirm that you wish to disable CarWings and asked to provide a reason why. Click 'Validate' when the appropriate option has been selected and you will get a confirmation message that CarWings has been disabled. You should also receive a confirmation via email. Once Nissan have resolved this issue it should be safe to re-enable your CarWings account and resume using features associated with it. Simply login to your account and follow the prompts on screen.

<u>Nissan's "CarWings"</u> is their telematics service and it's accessible in the UK via this page on their website (other countries will have different URLs).

Existing public domain knowledge

One of the key factors in publishing this now is the existence of multiple other public discussions about the unauthenticated API. The fact that only the VIN is required to invoke these services has been covered at length and published in locations including:

- A GitHub repository documenting the API including the observation that "All other operations take the DCMID and the VIN of your vehicle as parameters for authorizing the requested operation" (although the DCMID value is not actually required and is empty in many of the examples above)
- Another GitHub repository, this time a Python script to connect to and manage vehicle features via the API (also includes region codes for managing vehicles in other parts of the world)
- Yet another GitHub repository built to target an earlier generation of the service and referenced as inspiration for the previously mentioned project
- A blog post on reverse engineering the API which observes that "curiously, it seems like you just need the constant DCMID and VIN fields" (again, the DCMID parameter wasn't actually used in our tests)
- A forum post on integrating the data into <u>Domoticz</u> (a home automation system) which makes this observation: "No other authentication necessary!"

Whilst I haven't linked directly to the resources, they're easily discoverable via Google and demonstrate that there is ongoing public discussion via multiple channels, each documenting the lack of authorisation on the services.

Disclosure timeline

I made multiple attempts over more than a month to get Nissan to resolve this and it was only after the Canadian email and French forum posts came to light that I eventually advised them I'd be publishing this post. Here's the timeline (dates are Australian Eastern Standard time):

23 Jan: Full details of the findings sent and acknowledged by Nissan Information Security Threat Intelligence in the U.S.A.

- 1.30 Jan: Phone call with Nissan to fully explain how the risk was discovered and the potential ramifications followed up by an email with further details
- 2.12 Feb: Sent an email to ask about progress and offer further support to which I was advised "We're making progress toward a solution"
- 3.20 Feb: Sent details as provided by the Canadian owner (including a link to the discussion of the risk in the public forum) and advised I'd be publishing this blog post "later next week"
- 4.24 Feb: This blog published, 4 weeks and 4 days after first disclosure

All in all, I sent ten emails (there was some to-and-fro) and had one phone call. This morning I did hear back with a request to wait "a few weeks" before publishing, but given the extensive online discussions in public forums and the more than one-month lead time there'd already been, I advised I'd be publishing later that night and have not heard back since. I also invited Nissan to make any comments they'd like to include in this post when I contacted them on 20 Feb or provide any feedback on why they might not consider this a risk. However, there was nothing to that effect when I heard back from them earlier today, but I'll gladly add an update later on if they'd like to contribute.

I do want to make it clear though that especially in the earlier discussions, Nissan handled this really well. It was easy to get in touch with the right people quickly and they made the time to talk and understand the issue. They were receptive and whilst I obviously would have liked to see this rectified quickly, compared to most ethical disclosure experiences security researches have, Nissan was exemplary.

The ethics of discovery and disclosure

Just one last thing on how these vulnerabilities are discovered and reported because the ethics of this often comes up in my workshops. Risks like the one above were discovered by doing nothing more than using the app as it was intended to be used and observing the traffic going backwards and forwards. This is the mobile equivalent of opening your browser's dev tools and watching the network tab. Sometimes (such as with the realestate.com.au vulnerability I reported last year), this is all that's required. Other times and as was the case with the LEAF, it meant testing that the theory of one user being able to access another user's resource could be proven. In a situation where it's a car involved, you can't exactly head out and buy a second one in order to prove that when accessing one you can change a parameter to access another and whilst the proof above did involve checking the battery status of another vehicle, it didn't involve accessing any personally identifiable information or disadvantaging anyone in any way.

To me, it's this simple: if the intent is ethical and any findings are reported privately and immediately the moment you're confident a serious risk is present – and especially if it can be done without viewing anyone else's private data – then I'm comfortable that's in everybody's best interests. If you report before being confident there's a risk you end up wasting people's time and if you don't report, then you end up leaving people – and the organisation involved – at risk. A post such as this one is reviewed dozens of times over by myself and

where possible, a peer or peers (Scott, in this case) to ensure fairness and accuracy.

Summary

Nissan need to fix this. It's a different class of vulnerability to the Charlie Miller and Chris Valasek Jeep hacking shenanigans of last year, but in both good and bad ways. Good in that it doesn't impact the driving controls of the vehicle, yet bad in that the ease of gaining access to vehicle controls in this fashion doesn't get much easier – it's profoundly trivial. As car manufacturers rush towards joining in on the "internet of things" craze, security cannot be an afterthought nor something we're told they take seriously after realising that they didn't take it seriously enough in the first place. Imagine getting it as wrong as Nissan has for something like Volvo's "digital key" initiative where you unlock your car with your phone.

By pure coincidence, <u>this week Nissan unveiled a revised LEAF at the GSMA Mobile World Congress</u>. Clearly, like many car makers, their future involves a strong push for greater connectivity in their vehicles:

In a fully connected, fully mobile world, in-vehicle connectivity is an absolute must for today's drivers. That is why Nissan is proud to be at the forefront of developing efficient and reliable in-vehicle connected technologies that are available and accessible to all

Amongst the list of features the media release talks about being added to the NissanConnect app is the ability to remotely show the vehicle position on a map and analyse your driving. Whilst there are obvious upsides to drivers having access to these features, seeing them presented within the security implementation of the current app would be very worrying for obvious reasons.

I would have preferred to see faster action from Nissan. In my view, this is the

sort of flaw that needs to have the service pulled until it can be fixed properly and restored; it's not a critical feature of the vehicle yet it has the potential to impact its physical function and there's the privacy risk as well. Plus of course it's already being discussed publicly via that Canadian forum so the risk is well and truly out in the public domain already. I want to see Nissan secure this; I own a Nissan myself (albeit not a connected one) which I'm passionate about and am very invested in the brand, both emotionally and financially. But they do need to take action on this because clearly the current state is not satisfactory.

Update 1, 25 Feb, 12:00: Nissan has now taken the service offline.

Update 2, 25 Feb, 14:20: Per the <u>comment below</u> and further correspondence I've had via email, it appears that Canadian resources are still accessible using only the VIN.

Comments

I own a LEAF in the USA. Two things that make this hack a little less dangerous for owners:

- 1) the Leaf will only run climate control for 15 minutes if it's unplugged so little risk of draining the battery down unless someone is malicious enough to repeat this procedure every \sim 15 mins.
- 2) Nissan will text or email you a confirmation if your car's climate control is turned on remotely, or charging is stopped.

Not that the vulnerability isn't serious. But these do limit its potential for destructive use a little.

Troy: Problem is though, you can just turn it back on after it goes off. My understanding from Scott is that the notification is opt-in as well.

Scott Helme: This is correct and there are several ways to tackle this. The

request to turn the AC on is 'fire and forget', I simply send the request no matter what the current AC status is and it will either turn on if it's off, or have no effect if it's on. With this, I can simply send the request as frequently as I like, say every minute. That way, even if you turn it off, a maximum of 1 minute later, it will be back on. The AC on notifications are, as Troy says, optional and default off in all of the cars I've looked at. Knowing that your AC has been turned on doesn't help much given what I mentioned though. You'd just get an AC on email every 60 seconds!

__

This is about the 3rd time in a month that you've made the BBC, awesome! Excellent article, as always.

It's disheartening that this happens so often, and gets mentioned here a lot, and yet nothing changes, and the companies who make such mistakes are those who really ought to have known better.

Do those in charge of technology at Nissan, at banks and supermarkets and the countless others you've mentioned, as have others, never notice that this happens a lot, and wonder "gee, I should take a look at our own systems, make sure we haven't completed screwed up like all those other people..."?

I'm not a technology profession (I'm neither in tech nor a professional at all, but a physics student) and if even to me it's blindingly obvious that security matters, and what far too many organisations call security doesn't even pretend to be that for anyone with half a brain, then it really ought to be to those whose entire jobs it is to make sure things like this don't happen.

Troy: That exact statement is what really struck me after the Vtech breach last year – doesn't someone in their C-suite ever say "Hey, with all this hacking that's going on, we should check to see if everything is solid on our end". I've resigned myself to the fact that this is just going to continue and a week from now it will be another story with another organization.

C-level perception that cost of litigation < cost of mitigation, perhaps?

Troy: Possibly, although that implies a conscious decision to design weak security as it justifies the ROI. I don't believe that's the case with Nissan and the equation is also changing with the likes of the EU imposing fines of up to 4% of gross revenue in cases of gross negligence. Make litigation > cost of mitigation and things may genuinely change.

Epilogue

I'll be blunt - Nissan handled this incompetently. I wasn't that blunt when I wrote the blog post because Nissan seemed like the kind of company with enough lawyers to make my life painful if they took a dislike to the way I wrote the story up. The incompetence started with the very first phone call I had with them shortly after getting in touch with them via email. The security lead on the line couldn't wrap his head around how we were able to intercept the traffic: "but.. it's encrypted over HTTPS!" And thus began a lesson in active traffic interception on a device you have control of, the same lesson I'd given the guy that found this vulnerability after not more than an hour of training on the topic.

The next example of incompetence is the delays, and that much would have come across in the blog. I really tried to get them to do the right thing with the bug and fix it before going public, but they just wouldn't have a bar of it. They literally stopped replying to my emails right up until the time I said I was going public with it a month after discussions began. They asked me not to, but I did it anyway. I figured that by that time and given the number of opportunities I'd given them, their ability to get legal with me was pretty limited, but I wouldn't throw them too far under the bus in the blog post just to hedge my bets a little.

But that's not where the incompetence ended, not by a long shot. The service went offline pretty quickly after I posted the blog and it stayed down for many weeks, about 6 in total from memory. When it came back up, Scott pinged me privately:

"Hey, notice anything weird about this screen cap from the updated app?"

I looked carefully and whilst it wasn't immediately visible on first glance, I soon saw the problem. Down the bottom of the screen where access to location services is set, sat the following line of text:

"App explanation: The sprit of stack overflow is coders helping coders"

Whoa! WTF?! We quickly concluded that a Nissan developer had copied and pasted the text from Stack Overflow, obviously being completely oblivious to what it actually did. That was alarming in and of itself because remember, we're talking about a massive auto manufacturer writing code to control moving parts within a car and they're doing shit like this. But then we found the Stack Overflow post and realised we hadn't hit the bottom of the incompetence well yet. You see, that quote just above this paragraph is written precisely as it appeared in the app, all the way down to the mistyping of "spirit". It's spelled correctly on Stack Overflow so they can't have copied and pasted it, rather someone must have literally typed out every single character and still not understood what they were doing!

One other fun thing happened after the blog post went out and the service went offline. A guy called Melvyn Burchell posted a review of the app to the Google Play store that began as follows:

"Thanks to the actions of publication from a rather stupid researchers (Troy Hunt) nobody (at least in the UK) has any Carwings / Connect EV functionality anymore as Nissan turned it off as soon as the vulnerabilities received widespread publicity"

This review made me so happy that I literally printed it out and framed it as a little reminder that my good deeds don't go unnoticed \bigcirc

HERE'S HOW I VERIFY DATA BREACHES

A couple of years into running Have I Been Pwned and I was becoming increasingly conscious of a simple yet critically important challenge to the ongoing operation of the service: regardless of how good my intentions were, there was no escaping the fact that HIBP only exists due to a whole bunch of illegal activity. Most of the data in there came about as a result of criminal acts, some of which even landed the perpetrators in jail (the bloke who hacked LinkedIn and Dropbox is a perfect example). I honestly didn't know where I stood legally with this; could I hold the data? What would happen if I misattributed the source of the breach? Could either a company that had actually been breached or an individual within the breach go to town on me? Not only did I not know the answers to any of these questions, there were a lot of differing opinions from not just the general masses, but even those with legal backgrounds. Compounding the whole problem was the fact that the answers would likely differ significantly based on which jurisdiction you're talking about, but the internet (almost) knows no geographical boundaries. Data breaches definitely know no geographical boundaries!

I wanted to write this post to lay out where my own moral compass was pointing in terms of how I was handling breaches. I figured that if I had to live in the grey in order to run the service, I was going to do everything I could to position HIBP at the most legitimate end of the spectrum as I possibly could. I'd later write many more posts for precisely the same reason (or at least that would be a large part of the reason), and as time passed, I'd continually try to push HIBP out of the shadowy realms of data breaches into mainstream acceptance.

But I also wanted to write this post due to the frustration I had with people peddling clearly fake data breaches. I'd see it time and time again where someone would pop up and say "hey, Twitter had a massive breach, here's the

data". Now I don't know if they fabricated the story or the person who gave them the data fabricated it or someone else entirely different upstream did, it doesn't matter, what matters is that there was no Twitter breach. This would happen over and over again (even to this day), and it annoyed the hell out of me as it essentially boiled down to "fake news", a term which was gaining a lot of traction in 2016 when I wrote this post, albeit for different reasons.

07 MAY 2016



FRANKFURT | BY ERIC AUCHARD

Other headlines went on to suggest that <u>you need to change your password right</u> <u>now</u> if you're using the likes of Hotmail or Gmail, among others. The strong implication across the stories I've read is that these mail providers have been hacked and now there's a mega-list of stolen accounts floating around the webs.

The chances of this data actually coming from these service providers is near zero. I say this because firstly, there's a *very* small chance that providers of this calibre would lose the data, secondly because if they did then we'd be looking at very strong cryptographically hashed passwords which would be near useless (Google isn't sitting them around in plain text or MD5) and thirdly, because I see data like this which can't be *accurately* attributed back to a source all the time.

That's all I want to say on that particular headline for now, instead I'd like to focus on how I verify data breaches and ensure that when reporters cover them, they report accurately and in a way that doesn't perpetuate FUD. Here's how I verify data breaches.

Sources and the importance of verification

I come across breaches via a few different channels. Sometimes it's a data set that's broadly distributed publicly after a major incident such as the Ashley Madison attack, other times people who have the data themselves (often because they're trading it) provide it to me directly and increasingly, it comes via reporters who've been handed the data from those who've hacked it.

I don't trust any of it. Regardless of where it's come from or how confident I "feel" about the integrity of the data, everything gets verified. Here's a perfect example of why: I recently wrote about <u>How your data is collected and commoditised via "free" online services</u> which was about how I'd been handed over 80 million accounts allegedly from a site called Instant Checkmate. I could have easily taken that data, loaded it into <u>Have I been pwned</u> (HIBP), perhaps pinged a few reporters on it then gone on my way. But think about the ramifications of that...

Firstly, Instant Checkmate would have been completely blindsided by the story. Nobody would have reached out to them before the news hit and the first they'd know of them being "hacked" is either the news headlines or HIBP subscribers beating down their door wanting answers. Secondly, it could have had a seriously detrimental effect on their business; what would those headlines do to customer confidence? But thirdly, it would have also made me look foolish as the breach wasn't from Instant Checkmate - bits of it possibly came from there but I couldn't verify that with any confidence so I wasn't going to be making that claim.

This week, as the news I mentioned in the intro was breaking, I spent a great deal of time verifying another two incidents, one fake and one legitimate. Let me talk about how I did that and ultimately reached those conclusions about authenticity.

Breach structure

Let's start with an incident that has been covered in a story just today titled <u>One of the biggest hacks happened last year, but nobody noticed</u>. When Zack (the ZDNet reporter) came to me with the data, it was being represented as coming from <u>Zoosk</u>, an online dating site. We've seen a bunch of relationship-orientated sites recently hacked *and that I've successfully verified* (such as Mate1.com and Beautiful People) so the concept of Zoosk being breached sounded feasible, but had to be emphatically verified.

The first thing I did was look at the data which appears like this:

```
Windows PowerShell
                                                                                                      П
                                                                                                              X
         7@hotmail.com:wizemita
_valenza1955@libero.it:070755
gifmnamtrmnedfwpctxs
         alvez2@gmail.com:antequerana
            agbaba__36@hotmail.com:xuryniry
         piz@libero.it:camel
des@hotmail.com:33822734
an_1919@hotmail.com:gygomalo
         .bergrath@gmx.net:xozynewa
0401@web.de:rulatuji
         n_94@mail.ru:qaxojezy
         d_taieb75@yahoo.fr:sabah
         obinson@hotmail.co.uk:kakykipy
         da_@hotmail.com:\N
         day@hotmail.com:35027446
dibus@hotmail.com:benjamin
mellon@msn.com:mipyqofu
laly@hotmail.fr:didou1980
ra_kickboxing@hotmail.com:nokia3650
         @hotmail.com:28315514
         astankova@seznam.cz:\N
         otito@gmail.com:valyxaki
```

There were 57,554,881 rows of this structure; an email address and a plain text password delimited by a colon. This was *possibly* a data breach of Zoosk, but right off the bat, only having email and password makes it very hard to verify. These could be from anywhere which isn't to say that some wouldn't work on Zoosk, but they could be aggregated from various sources and then simply tested against Zoosk.

One thing that's enormously important when doing verification is the ability to

provide the organisation that's allegedly been hacked with a "proof". Compare that Zoosk data (I'll refer to it as "Zoosk data" even though ultimately I disprove this), to this one:

This data was allegedly from <u>fling.com</u> (you probably don't want to go there if you're at work...) and it relates to this story that just hit today: <u>Another Day, Another Hack: Passwords and Sexual Desires for Dating Site 'Fling'</u>. Joseph (the reporter on that piece) came to me with the data earlier in the week and as with Zack's 57 million record "Zoosk" breach, I went through the same verification process. But look at how different this data is - it's complete. Not only does this give me a much higher degree of confidence it's legit, it meant that Joseph could send Fling segments of the data which *they* could independently verify. Zoosk could easily be fabricated, but Fling could look at the info in that file and have absolute certainty that it came from their system. You can't fabricate internal identifiers and time stamps and not be caught out as a fraud when they're compared to an internal system.

Here's the full column headings for Fling:

CREATE TABLE 'user' ('duid' int(10) unsigned NOT NULL AUTO_INCREMENT, 'username' varchar(64) NOT NULL, 'password' varchar(32) NOT NULL, 'email' varchar(255) NOT NULL, 'email_validated' enum('N','Y') NOT NULL DEFAULT 'N', 'accept_email' enum('N','Y') NOT NULL DEFAULT 'Y', 'md5' varchar(32) NOT NULL, 'enum('N','Y') NOT NULL DEFAULT 'Y', 'md5' varchar(32) NOT NULL, 'm e m b e r s h i p 'enum('FREE','PROMO','GRANDFATHERED','BRONZE','SILVER','GOLD','ADMIN') NOT NULL DEFAULT 'FREE', 'join_date' datetime NOT NULL, 'birth_date' date NOT

NULL, 'location_id' varchar(8) NOT NULL, 'gender' enum('COUPLE', 'MAN', 'WOMAN', 'TS', 'UNSPECIFIED') NOT NULL DEFAULT 'UNSPECIFIED', 'seeking' set('COUPLE', 'MAN', 'WOMAN', 'TS', 'UNSPECIFIED') NOT NULL DEFAULT 'UNSPECIFIED', 'interested_in' set('FETISH','GROUPSEX','SEXUAL RELATIONS', 'ONLINE FLIRTING', 'OTHER', 'UNSPECIFIED') NOT NULL DEFAULT 'UNSPECIFIED', 'last_login' datetime NOT NULL, 'mobile_user' enum('N','Y') NOT NULL DEFAULT 'N', 'mobile_phone_no' varchar(16) DEFAULT NULL, `mobile_carrier` varchar(20) DEFAULT NULL, `discreet_profile` enum('N','Y') NOT NULL DEFAULT 'N', 'featured_profile' enum('N','Y') NOT NULL DEFAULT 'N', 'power_user' enum('N','Y') NOT NULL DEFAULT 'N', 'account_status' enum('ACTIVE', 'USER_DISABLED', 'ADMIN_DISABLED', 'SCAMMER_DISABLED') NOT NULL DEFAULT 'ACTIVE', `advert_id` varchar(25) DEFAULT NULL, `ip_address` varchar(16) NOT NULL, 'mtime' timestamp NOT NULL DEFAULT CURRENT_TIMESTAMP ON UPDATE CURRENT_TIMESTAMP, PRIMARY KEY ('duid'), UNIQUE KEY 'username' ('username'), UNIQUE KEY 'email' ('email'), KEY 'location_id' ('location_id'), KEY 'md5' ('md5'), KEY 'join_date' ('join_date'), KEY 'ip_address' ('ip_address'), KEY 'password' ('password'), CONSTRAINT 'user_ibfk_1' FOREIGN KEY ('location_id') REFERENCES 'geo_location' ('location_id') ON UPDATE CASCADE) ENGINE=InnoDB AUTO_INCREMENT=64192949 DEFAULT CHARSET=utf8;

The other thing in terms of structure is that the Fling data begins with this:

```
-- MySQL dump 10.11
-- Host: 192.168.1.28 Database: fling
-- Server version 5.1.41-enterprise-gpl-advanced-log
```

It's a <u>mysqldump</u> of the data with enough version and host info to again, create a much higher degree of confidence in the data not just for me in terms of how it "feels", but for Fling themselves to be able to verify.

I'm *very* suspicious of data presented in the way the Zoosk breach was and compared to Fling, you can see how both would impact my confidence levels in different ways. Let's move on though and increase that confidence level a bit.

Enumeration

Most websites will tell you if an email address exists on the site, you just need to ask. For example, enter an email address into Adult Friend Finder's password reset feature and they'll tell you very clearly if it's already in their database or not. It's not always that explicit, Ashley Madison used to disclose account existing by returning slightly different responses. If a site isn't facilitating enumeration on the password reset, then it frequently is on the registration feature ("this email address is already registered") and it's rare *not* to be able to simply plug in an email address and be told via one channel or another if it already exists on the site.

Enumeration risks such as these are not "silent" in that something like a password reset will send an email to the recipient. Whilst it's by no means compromising their personal security in any way, I also don't particularly want to inconvenience people. But there's a way around that and it provides another upside too.

Mailinator accounts in data breaches

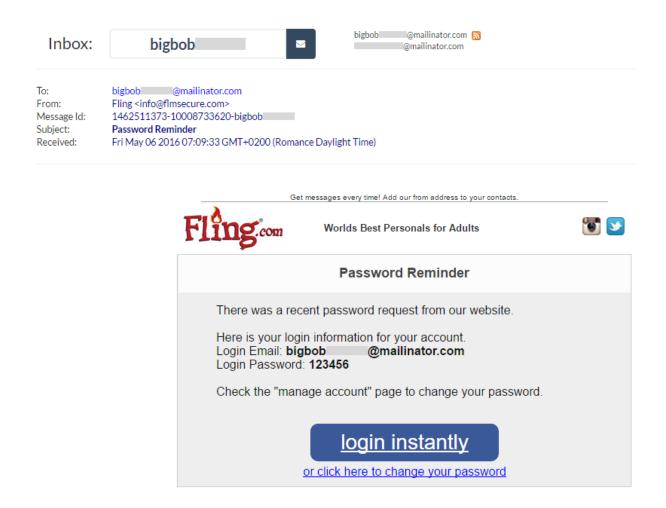
If you haven't used <u>Mailinator</u> before, you're missing out. It's an awesome way of standing up free, disposable email addresses and you can simply send a mail to <code>[anything]@mailinator.com</code> then check it on their site. There's also zero security and consequently, zero privacy. People often use Mailinator accounts simply as a means of passing the "please verify your email address" test that many sites pose before you can access them.

Mailinator accounts are perfect for testing enumeration risks. For example, the email address bigbob******@mailinator.com is the first one in Fling and if you plug that into their password reset form, you get this:

Forgot Password

Your password has been emailed to you.

Curiously, Fling returns exactly the same message when the email is entirely fabricated; fat-finger the keyboard and you'll get the same response. In that regard, password reset may not be an enumeration vector on Fling but it doesn't matter because when testing a Mailinator account, the reset email is publicly accessible anyway:



It turns out that Big Bob also has a password of commensurate security to his choice of mail provider, and this gives us another verification data point:

```
(282761, 'bbhaf', 123456', 'bigbob @mailinator.com', 'N', 'Y', 'Y
```

Of course you can only do this with a breach where the site actually emails the password which (fortunately) isn't that common, but you can see how each of these processes starts to build confidence in the authenticity of the breach. That can be confidence that it *is* genuine as well as confidence that it *isn't*.

The Zoosk data had way too many accounts that weren't checking out. *Some* Mailinator accounts would cause their password reset to respond confirming an email had been sent but many others didn't. It's possible that accounts had been deleted from their end post-breach (sometimes this is just a "soft" delete - the record is still there but flagged as inactive), but the low hit-rate wasn't inspiring

much confidence.

But there's another avenue I have available that's proven *very* reliable, and that's HIBP subscribers.

Verifying with HIBP subscribers

I'm now approaching 400k *verified* subscribers to HIBP, that is they've gone to <u>the</u> <u>free notification service page</u>, entered their email address then *received* an email at that address and clicked on a verification link. These are people who have an interest in protecting their online identities and they want to know about it when an incident occurs that impacts them.

What I've been doing with breaches that are harder to verify *or* I that want to have a greater degree of confidence in, is temporarily loading the email addresses into the SQL database in HIBP which stores the notification users (this doesn't contain the accounts the service allows you to search, those are stored in Azure Table Storage), then running a query that gives me results like this:

	Email		VerificationDate
1		nail.com	2016-05-06 05:25:38.70
2		net	2016-05-06 05:22:26.63
3		nail.com	2016-05-06 04:53:36.56
4		∄gmail.com	2016-05-06 04:48:32.04
5		.com	2016-05-06 04:37:35.43
6)gmail.com	2016-05-06 01:55:29.12
7		ail.com	2016-05-06 01:46:09.04
8		au@gmail.com	2016-05-06 01:28:30.27
9		gmail.com	2016-05-06 01:04:58.57
10		mail.com	2016-05-06 00:55:15.22
11		olny.cz	2016-05-06 00:46:47.46
12		ey@gmail.com	2016-05-06 00:46:32.38
13		.com	2016-05-06 00:05:39.28
14		.com	2016-05-05 23:51:53.23

These are the most recently verified HIBP subscribers who appear in the Zoosk data or in other words, those who have a recent recollection of signing up to the service I run. I'll take 30 of those and send them an email such as this one:

Hi, I'm emailing you as someone who has recently subscribed to the service I run, "Have I been pwned?"

I'm after your support in helping to verify whether a data breach I've been handed is legitimate or not. It's one that I need to be absolutely confident it's not a fake before I load the data and people such as yourself receive notifications. This particular one is quite personal hence the extra due diligence.

If you're willing to assist, I'll send you further information on the incident and include a small snippet of your (allegedly) breached record, enough for you to verify if it's accurate. Is this something you're willing to help with?

I send this off with everyone BCC'd so inevitably a bunch of them go to spam

whilst others are ignored or simply not seen for quite a while hence why I email 30 people at a time. People who *do* respond are always willing to help so I send them back some segments of the data to verify, for example:

This relates to the website fling.com which an attacker has allegedly breached. Your email address is in there with the following attributes:

- 1. A password that begins with "[redacted]"
- 2. An IP address that belongs to [redacted] and places you in [redacted]
- 3. A join date in [month] [year]

Does this data seem legitimate? Other indicators suggest it's highly likely to be accurate and your confirmation would be enormously helpful.

I sent this exact message back to a number of HIBP subscribers in the Fling data set and all of them confirmed the data with responses such as this:

That is indeed accurate. Lovely plaintext password storage I see.

There's a risk that people merely respond in the affirmative to my questions regardless of whether the data is accurate or not. However firstly, I've already found them in the breach and reached out to them - it's already likely they're a member. Secondly, I rely on multiple positive responses from subscribers so we're now talking about people lying en masse which is much less likely than just one person with a confirmation bias. Finally, if I *really* feel even greater confidence is required, sometimes I'll ask *them* for a piece of data to confirm the breach, for example "what month were you born in".

The Fling data was emphatically confirmed. The Zoosk data was not, although *some* people gave responses indicating they'd previously signed up. Part of the problem with verifying Zoosk though is that there's just an email address and a password, both of which could conceivably have come from anywhere. Those who denied membership also denied they'd ever used the password which appeared next to their email address in the data that was provided to me so the

whole thing was looking shakier and shakier.

Zoosk wasn't looking legit, but I wanted to try and get to the bottom of it which called for more analysis. Here's what I did next.

Other verification patterns

In a case like Zoosk where I just can't explain the data, I'll often load the data into a local instance of SQL Server and do further analysis (I don't do this in Azure as I don't want to put other people's credentials up there in the cloud). For example, I'm interested in the distribution of email addresses across domains:

	Domain	Count
1	hotmail.com	17674787
2	yahoo.com	8725857
3	gmail.com	4615946
4	hotmail.fr	3615199
5	hotmail.it	2185868
6	libero.it	977693
7	aol.com	956734
8	yahoo.fr	945474
9	live.fr	911914
10	hotmail.co.uk	897480

See anything odd? Is Hotmail having a resurgence, perhaps? This is not an organic distribution of email service providers because Gmail should be way out in front, not at 50% of Hotmail. It's more significant than that too because rows 4, 5 and 10 are also Hotmail so we're talking 24 million accounts. It just doesn't smell right.

Then again, what does smell right is the distribution of email accounts by TLD:

	TLD	Count
1	com	35816701
2	fr	6299313
3	it	5311579
4	de	1627925
5	uk	1622577
6	net	982845
7	es	846934
8	ru	706166
9	pl	702466
10	CZ	549265

I was interested in whether there was an unexpected bias towards any one particular TLD, for example we'll often see a heap of .ru accounts. This would tell me something about the origin of the data but in this case, the spread was the kind of thing I'd expect of an international dating service.

Another way I sliced the data is by password which was feasible due to the plain text nature of them (although it could also be done with salt-less hashes as well). Here's what I found:

	Password	Count
1	\N	1783889
2	123456	469762
3	123456789	143524
4	12345	62031
5	000000	43294
6	12345678	41543
7	111111	37890
8	1234567	37770
9	Password	34324
10	azerty	29795
11	zoosk	27559
12	123123	24463
13	qwerty	23836
14	1234567890	20551
15	654321	19391
16	666666	17181
17	badoo	15056
18	iloveyou	13522
19	andrea	13378
20	juventus	13213

With passwords, I'm interested in whether there's either an obvious bias in the most common ones *or* a pattern that reinforces that they were indeed taken from the site in question. The most obvious anomaly in the passwords above is that first result; 1.7M passwords that are simply the escape character for a new line. Clearly this doesn't represent the source password so we have to consider other options. One, is that those 1.7M passwords were uncrackable; the individual that provided the data to Zack indicated that storage was originally MD5 and that he'd cracked a bunch of the passwords. However, this would represent a 97% success rate when considering there were 57M accounts and whilst not impossible, that feels way too high for a casual hacker, even with MD5. The passwords which do appear in the clear are all pretty simple which you'd expect,

but there's simply not enough diversity to represent a natural spread of passwords. That's a very "gut feel" observation, but with other oddities in the data set as well it seems feasible.

But then we have indicators that reinforce the premise that the data came from Zoosk, just look at the 11th most popular one - "zoosk". As much as that reinforces the Zoosk angle though, the 17th most popular password implicates an entirely different site - <u>Badoo</u>.

Badoo is another dating site so we're in the same realm of relationship sites getting hacked again. Not only does Badoo feature in the passwords, but there are 88k email addresses with the word "badoo" in them. That compares to only 6.4k email addresses with Zoosk in them.

While we're talking about passwords, there are 93k on them matching a pattern similar to this: "\$HEX[73c5826f6e65637a6e696b69]". That's a small portion of the 57M of them, but it's yet another anomaly which decreases my confidence in the data breach being what it was represented as - a straight out exploit of Zoosk.

Another really important step though is actually confirming a breach with the owner of the site that allegedly lost it. Let's delve into that.

Verifying with the site owner

Not only is the site owner in the best position to tell whether the breach is legit or not, it's also just simply the right thing to do. They deserve an early heads up if their asset has been accused of being hacked. However, this is by no means a foolproof way of getting to the bottom of the incident in terms of verification.

A perfect example of this is the Philippines Election Committee breach <u>I wrote</u> <u>about last month</u>. Even whilst acknowledging that their site had indeed been hacked (it's hard to deny this once you've had your site defaced!), they still

refused to confirm or deny the legitimacy of the data floating around the web even weeks after the event. This is not a hard job - it literally would have taken them hours at most to confirm that indeed, the data had come from their system.

One thing I'll often do for verification with the site owner is use journalists. Often this is because data breaches come via them in the first place, other times I'll reach out to them for support when data comes directly to me. The reason for this is that they're very well-practiced at getting responses from organisations. It can be notoriously hard to ethically report security incidents but when it's a journalist from a major international publication calling, organisations tend to sit up and listen. There are a small handful of journalists I often work with because I trust them to report ethically and honestly and that includes both Zack and Joseph who I mentioned earlier.

Both the breaches I've referred to throughout this post came in via journalists in the first place so they were already well-placed to contact the respective sites. In the case of Zoosk, they inspected the data and concluded what I had - it was unlikely to be a breach of their system:

None of the full user records in the sample data set was a direct match to a

Zoosk user

They also pointed out odd idiosyncrasies with the data that suggested a potential link to Badoo and that led Zack to contact them too. Per his ZDNet article, there might be something to it but certainly it was no smoking gun and ultimately both Zoosk and Badoo helped us confirm what we'd already suspected: the "breach" might have some unexplained patterns in it but it definitely wasn't an outright compromise of either site.

The Fling breach was different and Joseph got a very clear answer very quickly:

The person who the Fling.com domain is registered to confirmed the legitimacy of the sample data.

Well that was simple. It also confirmed what I was already quite confident of, but I want to impress how verification involved looking at the data in a number of different ways to ensure we were really confident that this was actually what it appeared to be before it made news headlines.

Testing credentials is not cool

Many people have asked me "why don't you just try to login with the credentials in the breach" and obviously this would be an easy test. But it would also be an invasion of privacy and depending on how you look it, potentially a violation of laws such as the US Computer Fraud and Abuse Act (CFAA). In fact it would clearly constitute "having knowingly accessed a computer without authorization or exceeding authorized access" and whilst I can't see myself going to jail for doing this with a couple of accounts, it wouldn't stand me in good light if I ever needed to explain myself.

Look, it'd be easy to fire up Tor and plug in a username and password for say, Fling, but that's stepping over an ethical boundary I just don't want to cross. Not only that, but I don't *need* to cross it; the verification channels I've already outlined are more than enough to be confident in the authenticity of the breach and logging into someone else's porn account is entirely unnecessary.

Summary

Before I'd even managed to finish writing this blog post, the excitement about the "breach" I mentioned in the opening of this blog post <u>had begun to come back down to earth</u>. So far down to earth in fact that we're potentially looking at only about one in every five and a half thousand accounts actually working on the site they allegedly belonged to:



Mail.Ru analyzed 57 mil of the 272 mil credentials found this week in alleged breach: 99.982% of those are "invalid" motherboard.vice.com/read/hacker-27...



A Hacker Is Selling 272 Million Email Logins, But There's N...

Why not all data breaches are created equal, and why not data leaks are actually breaches.

motherboard.vice.com

That's not just a fabricated breach, it's a very poor one at that as the hit rate you'd get from simply taking credentials from another breach and testing them against the victims' mail providers would yield a *significantly* higher success rate (more than 0.02% of people reuse their passwords). Not only was the press starting to question how legitimate the data actually was, they were getting statements from those implicated as having lost it in the first place. In fact, Mail.ru was pretty clear about how legitimate the data was:

none of the email and password combinations work

A

Breach verification can be laborious, time consuming work that frequently results in the incident not being newsworthy or HIBP-worthy but it's *important* work that should - no "must" - be done before there are news headlines making bold statements. Often these statements turn out to not only be false, but unnecessarily alarming and sometimes damaging to the organisation involved. Breach verification is important.

Comments

I assume you're familiar with how users tweak their email address with the +something thing. That is, an email address of "name@domain.com" might get entered as "name+something@domain.com", where the "something" is usually related to the site they are giving their email address to. Some sites don't permit a "+" character in the email address, of course.

In your experience, how rare is this behaviour by users? Is there consistency of frequency across different breaches? are there obvious patterns in the +something value, like how you noticed "zoosk" and "badoo" in the password?

Troy: I can tell you exactly how rare it is: 0.03% of users apply it. At least that was the case in the Adobe breach: https://haveibeenpwned.user...

Having said that, you're right in that it's another indicator of the source and I only need a few instances of it to increase my confidence in the breach authenticity. In that way, it's the same as people creating one-off email addresses with the name of the service in them.

Speaking of one-off email addresses, I have the habit of using <gmail username="">+SiteName@gmail.com or u.s.e.rnam.e@gmail.com if that doesn't work (or a combination of both). Will HIBP check these variations for me as well or would I have

to register them all one by one for them to be checked?

Troy: There are only 0.03% of people who use that pattern and there's a high barrier to entry to implement it in HIBP so you'll have to search individually for them.

Epilogue

Years later when I went through the HIBP merger and acquisition process, I poured a substantial amount of money into lawyers. Part of their remit was to define the legality of the service operating as it was, an activity that involved many meetings with people in suits. We went round and round looking at Australian privacy laws, GDPR (Europe), CCPA (California) and other acronyms I can't even remember. The end result was a simple statement that was trotted out over and over again during the process: "We invite bidders to form their own views on the legality of the service". Good lawyer speak there guys.

One point I made in this post that's really stood the test of time was about notifying the site owner and as I said back then, "it's also just simply the right thing to do". Several years after writing this post I saw a perfect example of why this was so important: I was in San Francisco meetings with a heap of companies as part of the HIBP sale. I got 2 alerts to my email address that it had appeared in a data breach, each one coming from a separate service similar to HIBP. I won't say which breach or which services as I don't want to publicly throw any of these parties under the bus, but they all occupied what I'd refer to as the more legitimate end of the data breach reporting services spectrum. Anyway, I did indeed have an account on this service and soon enough, someone sent me the entire dataset. Having already received multiple notices, I assumed that this breach had been reported and was known, so I loaded it into HIBP. I was rushing around at the most stressful time of my life,

and I remember looking for a news story to link to in the breach description and not being able to find one. I ended up just putting a link to the company's website, made it live in HIBP and sent out the breach notifications to subscribers. Next thing you know there's news headlines all over the place about how this service had been breached and the data is in HIBP. Awkward.

But wait - there's more! A few days later I'm at the Blackhat conference in Vegas and I bump into the CEO of one of the aforementioned companies that had sent me a breach notice. We'd met before and he's a good bloke, so we had a bit of a chat. I quizzed him about the incident, and it became pretty apparent that they hadn't done any disclosure whatsoever. Because of HIBP's profile (and, I assume, significantly larger user base given it's all free), it was my service that raised public eyebrows even though other services had notified their own members. I left that conversation thinking firstly how I had a lot more responsibility to do disclosure right because of what HIBP had become, and secondly how as an industry we needed to do better. Much better. I have a pipedream of how to approach this, it's sitting in there on my "to do" list somewhere...

HERE'S HOW I HANDLE ONLINE ABUSE

By October 2016 when I wrote this blog post, my profile was growing and with that, the target on my back was getting bigger. If we work on the assumption that a fixed percentage of anyone's follower count are abusive lunatics, as the following increases then so does the abuse. But I suspect it's more than that; with success comes envy and resentment and it seems to exponentially increase abusive behaviour over and above what the mere number of followers suggests.

It's sad that I felt it necessary to write this post at all. It's sad that in the subsequent years I've received much, much more abuse. But the post helped me establish a compass for dealing with it; it set forth how I'd decided to respond and for the most part, it's still consistent with how I deal with abuse today. It's not a fun topic but it's an important one because now more than ever, online life is real life, and we all need tools to deal with this kind of crap.

17 OCTOBER 2016

originally wrote this post earlier on in the year. I honestly can't remember what the abuse was that led to it and frankly, that's probably for the best as it allowed me to re-read this and ensure it comes across as general advice rather than a knee-jerk reaction to a specific unpleasant experience. Whilst the simple process of writing it helped me get the episode off my chest at the time, I've decided to post it now because I think it's important, both for others who encounter nasty behaviour online and for myself when I next do.

Unfortunately, if you spend enough time online and especially if you're public enough, this is something you're going to have to deal with sooner or later.

Here's how I handle it.

Abuse

I'm writing this outside the context of any recent events for reasons that will become clearer as you read on, but after the last abuse incident I thought I'd finally jot some things down. Mostly this serves as a reference point – something I may direct people to in the future – but I also write many of my blog posts as a way of forcing me to think clearly about a topic and articulate it in a cohesive fashion.

It may not be something that many of you would have expected, but I've often found myself at the receiving end of online abuse. As time goes by and I get more exposure or profile or whatever you want to call it that puts me in front of more people, I get more vitriol from online antagonists. Let me explain what I mean by that, the types of abuse I get and how I've elected to handle these incidents.

What I think constitutes abuse

Let me clear this up first because I appreciate there's a degree of subjectivity to all this. The sorts of online abuse I get ranges from minor name-calling to slurs about my competence or professionalism to serious threats related to my personal life (I've come close to contacting the police in the past). I'm not going to detail what any of these actually were here as I simply don't want to give the trolls the airtime (more on that later), but I do want to describe some of the broader behaviours.

What I don't consider abuse is vehement disagreement with my points of view, finding factual faults with things I've written or said that are incorrect or any

other sort of constructive argument that I may not agree with, but is aired without malice or spite. It's the stuff that's said first and foremost to insult or cause harm that I put in the abuse bucket. This is *particularly* true when it's done from behind the veil of anonymity.

Very frequently, this is aired publicly via Twitter, in blog comments on troyhunt.com or via other online channels. Only very occasionally does it come via private means and it has *never* come verbally either face to face or via the phone. At times where I have actually engaged with the other party and offered to talk to them, the opportunity has *never* been taken up.

I should also be very clear that this is nothing like the abuse you hear of some people copping online; repeated threats to safety or family, prolonged "campaigns" of torment, racial or sexual abuse – all of that is a world apart from what I'm describing here. What I cop is merely nasty vitriol in comparison. In fact, very often it's the sort of thing I'm teaching *my six-year-old* is just inappropriate, nasty behaviour and I'm teaching him this because it's the sort of thing you expect from kids, not grown adults.

Let me explain some of the grievances that have come up multiple times before and I'm going to address them here once and for all.

I'm "profiting from security"

<u>The very first blog post I wrote</u> was in 2009. The first dollar of any significance I recall making out of security was when <u>my first Pluralsight course</u> went live four years later. There may have been some other inconsequential amounts but what I can say for sure is that until Pluralsight kicked in, 90% plus of my income came from working my arse off in a very corporatey role at Pfizer.

One thing that many people don't realise is that almost every time I talk at an event – including when I travel to the other side of the world to do it – I don't

earn a cent (there are a small handful of rare exceptions). Actually, I make negative money because a huge amount of time goes into not just the travel, but the preparation as well. Between conferences, podcasts and interviews, I've done hundreds of talks and almost never made a cent directly from them. These events are about meeting people and increasing my exposure, not just in terms of me putting my name out there, but me getting exposed to other really smart people. My experience has been that the best way to ultimately be personally successful in this area is to do as much as you can for free!

In more recent years, the work I've done has begun to pay well, almost entirely off the back of Pluralsight and the workshops I run. It pays well because it's in demand; there's a dearth of good security content targeted at developers and evidently the approach I take to explaining it is popular, something I make no apologies for. Which actually brings me to my next point: who my content is for.

I'm not explaining things "the right way"

Let me give you a perfect example of this: I've often seen disparaging comments about the use of the <u>Wifi Pineapple</u> to demonstrate security concepts. I'll see comments about how it's trivial or a "script kiddy" tool or how real men build their own devices and so on and so forth. What a lot of people seem to miss – and this predominantly comes from security professionals – is who I'm talking to.

The material I create, whether that be on blogs or at talks or in workshops, is very heavily biased towards software developers. Not only is that my background, but I believe that's where I can make the most difference to security; at the point where software is being written. In a case like the risks the Pineapple demonstrates, the vast majority of developers are unaware of how easily traffic can be hijacked or the risks behind practices such as loading login forms over HTTP. My goal is to make these concepts easily consumable to them

and the most impactful possible way I've found to do that is by showing how you can order a \$100 device off the web, pull it out of the box and 5 minutes later be hijacking traffic. That resonates more with that audience than rolling your own MitM tools ever will.

I fully appreciate that the way I'm explaining security to developers is not the way some security professionals would like to consume it themselves; it's not meant to be and the very fact that developers often get exposed to security in ways they have trouble consuming goes a long way to explaining why so many of them have such a poor grasp on it. In fact, that's the very reason I started getting involved in security many years ago – because of the friction I saw between developers and security teams.

There are people who understand many of the concepts I talk about at a greater depth than I do. Some of them are specialists in various niches, others have simply been focusing on specific things for longer. What I've found my strength to be is in explaining concepts in a way that's consumable by the people I speak to. I hope that makes sense and whilst not everyone will agree with the way I present some of these concepts, they can at least appreciate *why* I put them forward in that fashion.

"Tall poppy syndrome"

This is a term we hear a lot in Australia and whilst there might be different descriptions for it overseas, it generally means the same thing:

The tall poppy syndrome is a pejorative term primarily used in the United Kingdom, Australia, New Zealand, and other Anglosphere nations to describe a social phenomenon in which people of genuine merit are resented, attacked, cut down, or criticised because their talents or achievements elevate them above or distinguish them from their peers. This is similar to begrudgery, the resentment or envy of the success of a peer.

In other words, people being pissed because you've done well. I remember learning this term as a kid when you'd see someone getting cranky because someone else has just driven past in a nice car. I'm not sure if tall poppy syndrome is actually jealousy or just the view that someone else shouldn't be successful in what they're doing, but frequently this seems to be the undertone of abusive messages I receive.

Sometimes, the underlying resentment when a positive event occurs is particularly raw. I've seen cases where I've announced something or had some level of success or positive coverage and amongst the outpouring of absolutely awesome feedback, is one lone dissenting voice. Not a subtle disagreement, but outright vitriol. It's happened enough times in the past to be something I now expect, yet it never ceases to amaze me just how *opposite* that voice is to all the other ones.

Abuse like this doesn't have to be cogent or well-articulated and indeed the position of "I don't like you because you've achieved some level of success" is neither of these things. Yet somehow, antagonists taking this position seem to find time to commit to explaining how little attention others should be paying!

I'm a Microsoft / Lenovo / [anything else] shill

I'm certainly not alone in copping flak for affiliations and I can understand the *assumption* of me being incentivised to say positive things about companies that give me things, but there's a fundamental misunderstanding of the order in which these things occur. I'm a Microsoft Regional Director and MVP because I spent years writing about their technologies while receiving nothing from them. I'm a Lenovo Insider because I spent *decades* buying their gear and sharing my experiences publicly before they gave me a thing.

The irony of some of the abuse I get (and certainly some people do get very angry about my affiliations), is that I'll be reading about how I'm a Microsoft fanboy whilst using my iPhone (I don't want a Windows phone) or am beholden to Lenovo while reading *that* on the W540 I bought with my own hard-earned cash a couple of years ago. Independence and trustworthiness is *massively* important to me to the point where I push back on anything which has even an inkling of a chance of not being consistent with that. If it's not something that's an accurate reflection of my own independent views, I outright refuse and that's the end of it. It's that simple.

Funnily enough, I've often copped flak (I'll stop short of calling these incidents "abuse") about my ongoing promotion of tools like <u>Freedome VPN</u> and <u>1Password</u>. I've never received a cent from either of them and I've bought every single version of their respective products at retail prices out of my own wallet! I have no financial incentive, yet I influence people to purchase them *simply because they're very good!*

I recently spoke to someone in another position of influence with a similar affiliation to another large tech company and was *very* surprised at the pressure they had to not be seen with competitors' equipment. That's never the case with Microsoft or Lenovo and frankly, we're all that much better off that the opinions of those of us involved in their programs genuinely are independent, regardless of what those who like to hurl insults from the sidelines may think.

Actions I take when receiving abuse

I've changed my approach over the years as I've gone through various nasty experiences. Earlier on, I'd be tempted to confront antagonisers and challenge their negative perceptions – reason with them, if you like. Other times I've allowed followers to argue with them via channels such as Twitter and blog comments, sometimes I've even RT'd their ridiculous comments purely to invite

a torrent of defensive comments. These days, I'm trying to be much more passive.

One common thing among these individuals is that they want a fight. They're out there to argue and debate and do whatever they can to piss you off and consume your time. I now *mute* them at the first sign of the behaviours I described above. Twitter is easy because there's literally a mute feature and for anyone else who finds themselves in the same position, I highly recommend this. It's different to "blocking" them in that they can still see my timeline and as far as they know, I just haven't seen their message or I'm ignoring it – the joy of muting is that they don't know. Blocking is more "passive aggressive" and it's implicit engagement; IMHO, simply ignoring them from the outset is less confrontational. If it's comments on other blogs or social sites, I self-mute or in other words, I simply don't go back to that discussion. I make a conscious decision that doing so would be counterproductive and I simply tune out and go do something constructive.

Comments on my own blog are different, simply because that's *my* place and like others who run a blog, I get to decide what stays and what goes. After a nasty incident some years back, I created a page titled <u>Comments on troyhunt.com</u> which I link to just next to the comments section on each blog post. The bottom line is that if someone is abusive then I'll delete the comment and likely ban them. I've already clarified what I mean by abuse and in blog comments it's often insults or cheap shots without even an attempt to add something constructive to the discussion. I don't have any moderation before a comment goes live because I *want* people to come to my blog and discuss the content there, but when the goal of the comment is purely to antagonise without adding value to the content then that's it – it's gone.

When I look back at how I've handled previous incidents of online abuse, there are times where I wish I hadn't engaged. Perhaps the person was literally having the worst day of their life or had gone through a few too many glasses of the merlot or maybe they were just proverbially kicking the dog. There were

occasions where my engaging with them didn't work out well for either of us; for me because I wasted time debating with them when I could have been doing useful things, for them in various other ways which they likely now regret.

By pure coincidence, after writing this but before publishing, I read this about Robert Scoble:



This is just nasty. I'd stop short of calling it abusive, but it's the sort of behaviour that makes the guy look like a dick. No qualification of what it is about Robert he doesn't like, nothing constructive or insightful, just a nasty comment that

many people would find hurtful. That's not out of the ordinary, but it's Robert's response that really resonated:



Robert Scoble

I have some empathy. Maybe some serious shit is going on in his life and he feels like he needs to kick somebody. I am secure enough in my standing that if it helps him that's cool.

And this is precisely the point: there will be whingers who for no apparent reason just want to rant. No matter how well-regarded you become at what you do (or perhaps *because of it*), this stupid behaviour will appear and you can't help but feel a little bit sorry for the individual who resorts to it. I'm secure enough that I can happily ignore it and I'm not going to devote emotional energy to them which could be used to actually do good things.

Also, read both the cranky guy's comments and Robert's response – you actually come away from that with a greater respect for Scoble despite the original negative comment. In fact, for the vast majority of us, cranky guy has caused precisely the opposite effect to what he set out to achieve; he looks like a dick and his target comes out looking level-headed and having earned a new degree of respect from a bunch of people, myself included.

Here's a question to ask yourself if you recognise your own behaviour in any of this: would you willingly approach me face to face at a conference and say the same thing? Would you look me in the eye and repeat the abuse with the same conviction as you do – often anonymously – from behind the keyboard? If the answer is "no" then think about how invested you really are in your views and if perhaps it's something you shouldn't be saying in the first place.

Often these individuals are just exercising bravado that deserts them once they're away from either anonymity or the perceived invisibility that being on the other end of an internet connection gives them. Their better judgement and common decency is put aside in ways it simply wouldn't be were they not behind those veneers. But whilst they're behind the "protection" of an IP address and feeling as though they have no accountability, there's very little point in debating things; rational conversation is the last thing they're interested in.

It's literally a small fraction of 1% of people I interact with who decide to behave in this way and that's likely representative of most people at the receiving end of this sort of behaviour. So for me – and my advice to others as well – is that the right approach is unless it becomes an issue you simply can't avoid confronting, do your utmost to ignore it and move on. Angry or antagonistic people like an audience, better you don't give them one and they go elsewhere to find it.

The best defence: go and do awesome things!

There will always be cranky people who just want to get under your skin. We've no doubt all had that in the school yard before and many of us have had it in the workplace too. Online is a different story though and one of the best possible things you can do is drown out the negative noise with positive things.

I can't recall who I heard originally say it, but I distinctly recall a quote very similar to this:

You can't remove all negative things about you from the internet, the best thing you can do is to flood the web with positive things

And that's precisely what I intend to keep doing. In fact the abuse is motivation to go out and do great things that people love and want to share positive feedback about; more talks, more courses, more support for data breach victims via Have I been pwned – all of this makes the 99.x% of people I interact with on the web happy and that remaining fraction of a percent will simply need to accept that their abuse is being drowned out to the point where very often, I simply never even know it's occurred.

Comments

Reading this reminded me of the story of Curtis Woodhouse and what he did to one internet troll. Which is one extreme and different way of doing things.

When Curtis Woodhouse lost his English light-welterweight title on points to Shane Singleton on Friday night, he was branded a "disgrace" on the social networking site by 'Jimmyob88', who has reportedly been abusing Woodhouse on Twitter for months.

The boxer was so enraged with the tweets that he offered his followers a £1,000 reward if they could help him locate the culprit. Woodhouse's growing number of Twitter followers chipped in and managed to track down his troll.

Woodhouse set off to find his troll, tweeting a photograph of the street on which Jimmyob88 lived. "Right Jimbob, I'm here," he wrote, adding: "Someone tell me what number he lives at or do I have to knock on every door #itsshowtime."

Realising the error of his ways, Jimmyob88 replied: "I am sorry it's getting a bit out of hand. I am in the wrong. I accept that." A triumphant Woodhouse went home, joking that he could have saved himself some petrol money by blocking his cowardly abuser.

At the end of it he acted like every bully when confronted, and said it was only a joke.

Troy: Nice. More than anything, this shows how unwilling people are to back up their abuse in person. Even without the threat of violence, I'm yet to see an antagonist express the same vitriol face to face.

I had someone complain at me about profiting from talking at conferences. They just didn't get it that I don't get paid for it. Large conferences travel and hotels are covered sure, but for all the user groups I do, I pay my own train and hotel costs. I do it because I enjoy it.

Troy: More to the point, so what if you did? You know as well as I do how much

effort goes into preparing for the events we speak it and we usually get no short-term financial gain from it. But it's a long game that makes sense over time and it's worked out in very positive ways for me which I make absolutely zero apologies for!

I can't believe (figuratively) you're not paid for talks. don't people have to pay to go to them?

Troy: This surprises many people, but it's *extremely* rare for any of the speakers to be paid for major tech events. The norm is that travel and expenses are covered but your time isn't so whilst you get to attend a conference for "free", it means forgoing other paying work. In my case when I worked a full time job that meant usually attending events with my annual leave time. That's time which has a clear dollar value as I was paid out for remaining leave time when I left the company. Now as an independent, if I even do just a 1 hour talk overseas I'm pretty much taking a week off from earning money because everything is so far away from Australia, then you add jet lag and recovery and there goes 5 days.

Keynotes at large events are probably the only exception as are events outside the tech / developer conference scene. Conferences charge delegates to attend but many of those events are run by commercial entities that still need to turn a profit and they rely on speakers being available without charging for it. On the flip side, it increases your exposure and builds your "brand"; most of what I do commercially today is as a result of getting my name out there and a huge part of that is due to being visible at conferences. Many of the workshops I've consequently run and earned well from are due to people having seen me speak at events.

But yeah, when you see myself or others at events like I'm at today (and the one I'll be at tomorrow), remember that there's not a single cent coming from the organisers and going into my pocket.

In the US a significant percentage of conferences do not even cover T&E. Non-US conferences seem to be much better at paying for it.

Troy: I outright refuse to do those. It's one thing to not pay you for your time, it's quite another to also ask you to pay for all your costs in order to support their commercial venture.

"On the Internet, nobody knows you're a dog" comes to mind. Perhaps a more appropriate cartoon would be "On the internet, everyone's got an opinion".

Think you've got it bad? There are people running huge sites out there - youtubers - with millions of followers that take some serious abuse. And there's been plenty of casualties too - youtube channels and sites closing etc.

Dig a little deeper and speak to anyone else that has an online presence with decent exposure and you will find experiences to be strikingly similar - par for the course as they say.

I don't think anyone really thinks about the abuse part when they start a site or presence. For most of us, we start a site because we are passionate about something. So when negative and personal comments are directed at you, it's only natural to take it personal and act defensively.

Changing your outlook on things definitely makes a huge difference, but you still need to deal with incidents on a case by case basis and ultimately they still affect you.

In the end it is what it is : negative comments and abuse is part of having an online presence.

Troy: I'm very conscious that there are people out there that cop it much worse than me. Many of them are probably also more vulnerable (particularly younger people) and many of them are females who cop abuse of a sexual nature as well. I have somewhat of a luxury in being able to ignore people, it's much harder for some others I can totally get why they end up withdrawing altogether

Epilogue

This is one of those posts that has a hidden backstory, one that I didn't feel comfortable referring directly to in the post. A couple of years before writing this, I'd been pretty aggressively targeted by a small handful of people in the infosec community in Australia. I still have the messages as I got into the habit of saving anything that started to resemble a pattern of abusive behaviour in case I'd need it later on so I can relay precisely what was said here now. It boiled down to constantly accusing me of being "fake". I didn't know what I was doing in this industry. I had no experience. I was profiting off writing blog posts about poor security. And the final catalyst for the abuse? I'd written a blog post about transferring domains from GoDaddy to DNSimple. Yeah, go figure...

Anyway, the abuse became so constant that I dug a little deeper and discovered the person responsible worked for a mate of mine who ran a local penetration testing company. I gave him a call, we had a chat and... he fired the guy. Boom. Gone. It wasn't just his targeting of me, the guy had a history of misbehaving that didn't reflect well on his employer given he was publicly associated with the company. Around the same time, Twitter also suspended his account after constant harassment of a female infosec journalist. Classy guy. The last I ever heard of him was via an anonymous post to my blog which, with great certainty, I'm confident came from him. I've not redacted the language because to do so to the point where the words are obscured wouldn't leave much content so here it is in all it's original (all caps) glory:

Troy McShittyMcCuntBalls (Guest):

FUCK OFF TROY. YOU'RE JUST A CUNT WHO THINGS HE KNOWS ABOUT SECURITY.

STOP RIPPING OFF PEOPLES CONTENT AND FUCK OFF. NO ONE WANTS TO SEE YOUR

USELESS FACE AT CONFERENCES. STAY THE FUCK AT HOME AND ROT. ASSWIPE.

This is why much of the blog posts reads like me justifying my existence rather than just talking specifically about how I deal with abuse. In many of those headings I was flagging the topics people were abusing me about because I felt I needed to defend myself. It felt super shitty to write because I felt like I was under attack, and I certainly didn't have the thick skin I have now back then.

It took me a couple of years before I felt comfortable enough writing about online abuse and the example above was the first time I'd really copped it. I'd get a lot more in the years to come ranging from the sort of thing you just read to threatening behaviour about suing me for various things to outright death threats. But clearly, it didn't stop me, I just got better at handling it. These days I find myself muting people on Twitter very quickly which usually avoids the sort of escalation I would have previously been happy entertaining.

HERE'S EVERYTHING THAT GOES INTO A MASSIVE INTERNATIONAL SPEAKING TRIP

I didn't plan to create a 10,000+ word epic when I sat down to write this post but adding to it each day over the 3-week journey quickly became a routine that led to 21 short stories. My original thinking was to highlight the huge amounts of work that go into making a trip like this happen and I *think* I got that message across, but I also think that looking at it again now, it does look kinda glamorous.

Do read the actual *contents* of the post rather than just flicking through the pictures. There are so many instances of the word "lonely". So many references to being tired. That's not captured in the tweets and images because we tend to share the highlights of our lives, not the other way around. Read the text in this post because that tells the real story about the sacrifices that needed to be made to make it happen.

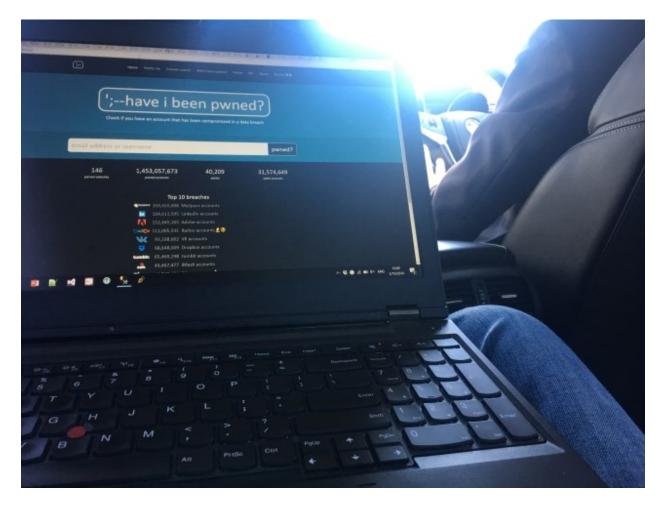
24 OCTOBER 2016

International travel can look pretty glamorous from the outside and certainly it has its moments. But what many people don't tend to see (and indeed what's less interesting to share in 140 char tweets), is just how arduous it can be. So instead of just showing the good bits, I thought I'd jot down a bit more about just how much stuff I fit into one of these trips, my fifth (and last) big international one for 2016. If you think it's all fun and games or if you're just curious about what on earth it is I do, read on and do keep reading too because whilst it'll all start out looking nice, it'll inevitably have some very hard and probably very dark moments.

Here it is, all the good bits and all the bad bits captured candidly as they happen:

Day 1, Sunday October 2: Leaving home

A car picks me up from home just after 9am. Good time of day because I get to spend the morning with the family and don't have to start out at a crazy hour. One of the neat things about flying on Qantas business class tickets is a pickup and drop-off service which makes quite a difference. (Last year I wrote about how I justify more expensive seats and as you'll see, I maximise the space to do productive things.) It also means that right from the outset, I can actually get some work done:



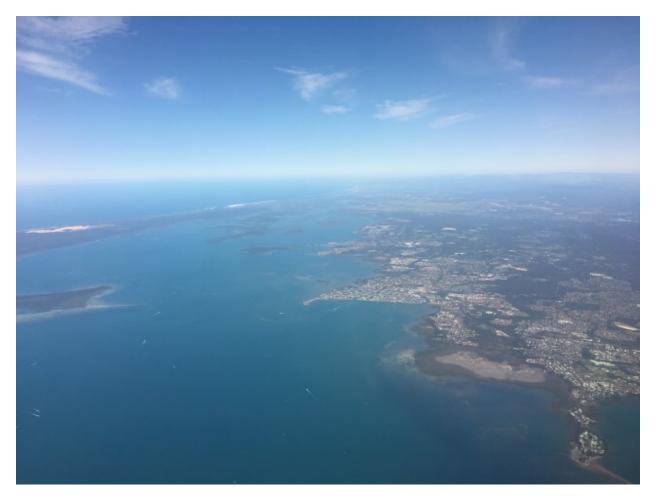
I get most of a new data breach loaded into <u>Have I been pwned</u> (HIBP) which is

pretty good use of the time. I try and max out every spare moment of travel I get and an hour here with internet connectivity is pretty useful.

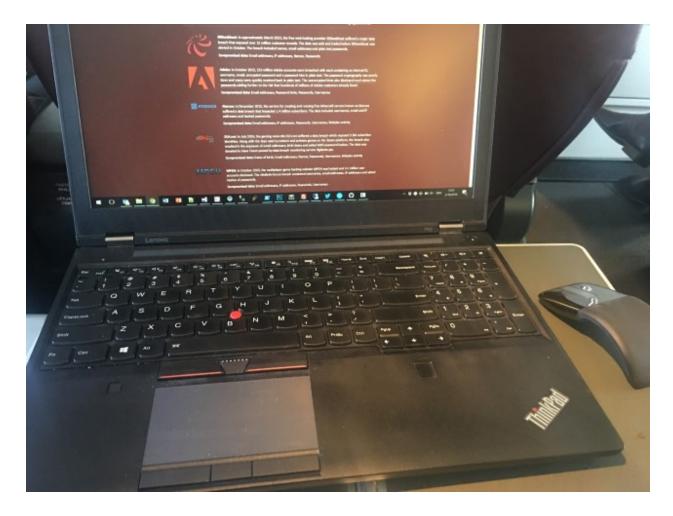
I take carry-on luggage only as it saves a heap of time at check ins and baggage claims not to mention all the walking I know I'll do with full kit in tow. I travel as light as possible, but "light" is a relative word:



And yes, everything in there I need either as a primary piece of equipment or as a backup. With all that and a carry-on bag for clothes, I can go from the car to sitting in the domestic lounge in 5 minutes tops. Domestic? Yeah, I have to fly from Brisbane to Sydney first so totally the wrong direction, but flights from Brisbane to Europe either weren't available on that day or were ridiculously expensive so here we are. I wrap the data load up in the lounge and jump on a domestic flight, flying back over home as I go:



It's a one-and-a-bit hour flight so I can't do much by the time I have some lunch too, but I get more done on another HIBP breach:



Get to Sydney and it's a bus to the international terminal then customs then another lounge. It's the last bit of wifi I'll get for a long time.







About to jump on the big plane to London, only 24 more hours to go!

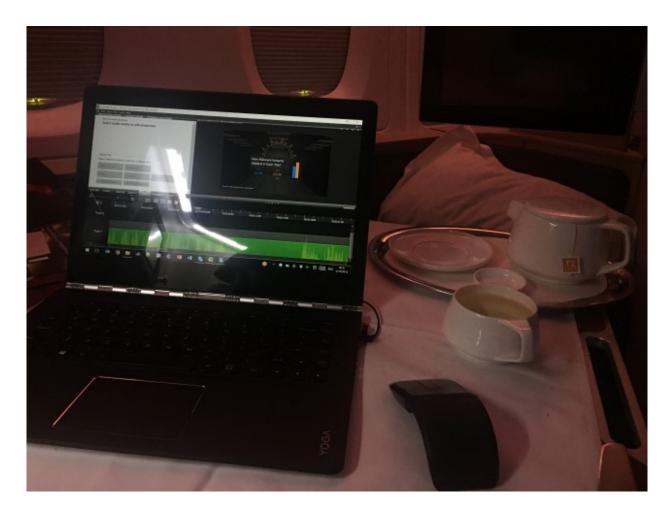
7:16 PM - Oct 1, 2016

○ 19 See Troy Hunt's other Tweets

Onto the big plane and it's A380 all the way which is nice. First class is also nice, but let me explain - I fly business class because it means I can work, sleep and make way more productive use of the time - there's a clear ROI for me. Lots of business class travel gives me lots of points... points I can't do anything with. Seriously, I've even stopped trying because I've *never* been able to use Qantas frequent flyer points to book a business trip where I need to leave and arrive on specific days with a few months' notice at most. I've even tried to book the whole family to Vancouver and back on economy (which is just fine for a family

holiday), and that's enormously restrictive too. Flexible on days a year in advance? Maybe, even some domestic flights I can use them on but other than that the only thing I've been able to do with the points is spend them on upgrades and even then, I'm accumulating them faster than I can spend them. I'll return from this trip with more than I left with courtesy of what I'll still earn from the business ticket.

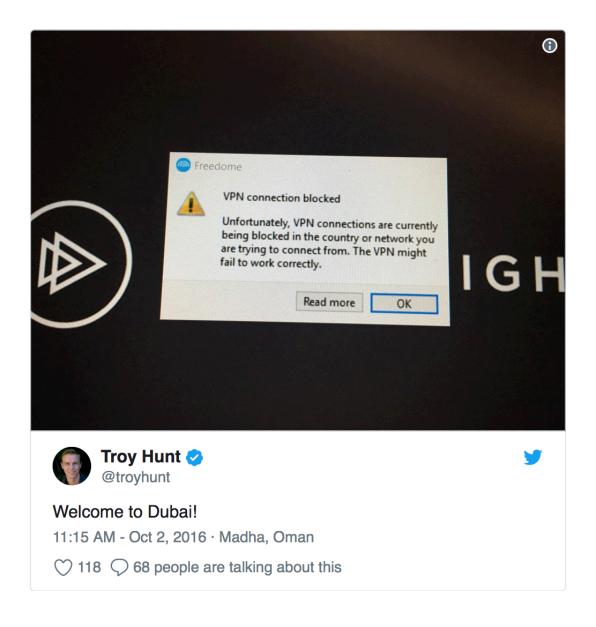
On the way to Dubai I get a bunch of coding done for a feature someone was after for HIBP plus get through a module and a half of editing for my next Pluralsight course. I always try and record a course before travelling as editing is a really good use of flight time (no internet, I'm bored anyway and editing is tedious). I take both my Lenovo P50 (I can't begin to tell you how much I love this machine!) and my Lenovo Yoga 900 because I've *never* been able to get enough power from plane sockets to charge them so I burn through the battery on one then roll over to the other:



I'm really careful on international trips to plan sleep: I figure out when I need to sleep to ensure I can acclimatise immediately at the other end. I take it easy on the alcohol and have herbal tea and a lot of fruit. I get about 8 hours of restless sleep which frankly, is pretty good.

Day 2, Monday October 3: Dubai to London

I get to Dubai after midnight and attempt to fire up Freedome VPN:



Oh yeah, the whole we're-blocking-vpns-so-you-can't-use-voip thing. It was fine when I went through only a few months earlier, but now it's no VPN for me. That means there's no way I'm RDP'ing into any important services or connecting via SSMS to HIBP (and yes, they have encrypted transport layers anyway but they're way too valuable to risk). There are various other ways around this, but with only an hour on the ground there's not much point.

The shorter Dubai to London leg gives me enough time to edit another Pluralsight module so that's almost half of the 6-module course now done. I get

another 2 and a half hours sleep then breeze through customs and baggage claim (another benefit of business travel with fast track tickets and only carry-on luggage). This is the ROI I speak of in paying for better seats: I've arrived rested and having been pretty productive, neither of which I can do with my 6'5" frame in small seats.

Jump in a waiting car then struggle though London traffic as a strange orb normally foreign to the UK rises above the horizon:



I'm at the hotel by 9am and they let me check in early. I shower, then head straight out for the day. I've planned a massive walk and some stops to meet people because frankly, it's the best way of acclimatising quickly. Full on work

starts the next day so the last thing I'm going to do is sleep at the wrong time for the new environment. Plus, it gives me a chance to head to <u>Hamleys</u> and buy my son something for his birthday, an event which I'll miss on this trip. I walk through the store with him on Facetime video and let him choose something that'll actually fit in my baggage (IMHO, a damn cool idea because Hamleys is awesome!)

Drone now in hand, I go for a walk to Hyde Park:

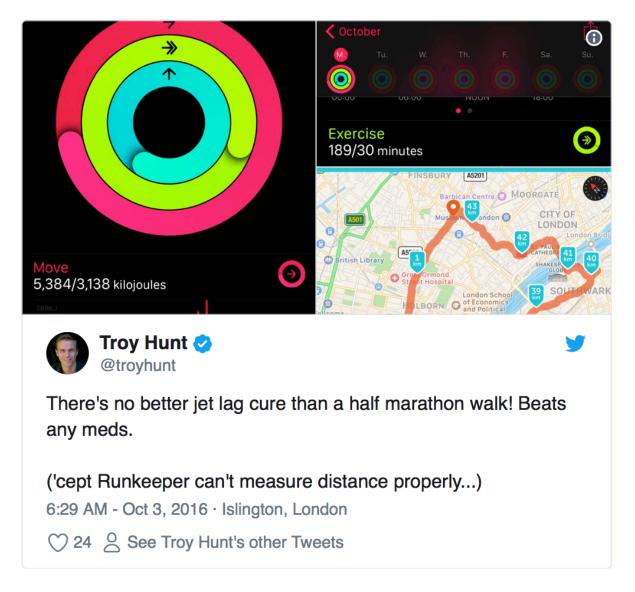


Then meet up with a company I've been talking to about various bits and pieces, have a coffee and a tour then some lunch. More walking, then coffee in the arvo with some Twitter contacts:



This is actually really cool doing impromptu stuff like this and meeting new people, something I really recommend if you're travelling somewhere new and have downtime. I'd been talking online with the Twitterers who came along and I've never had a bad experience catching up IRL with people like this.

By the time it's all done, I've had a good walk:



I eat at a normal time, go to bed just a little bit earlier than usual then sleep for 10 hours and get up at a normal time. This is invaluable - body clock is good! Anything remotely glamorous about travel ends tomorrow.

Day 3, Tuesday October 4: London workshop day 1

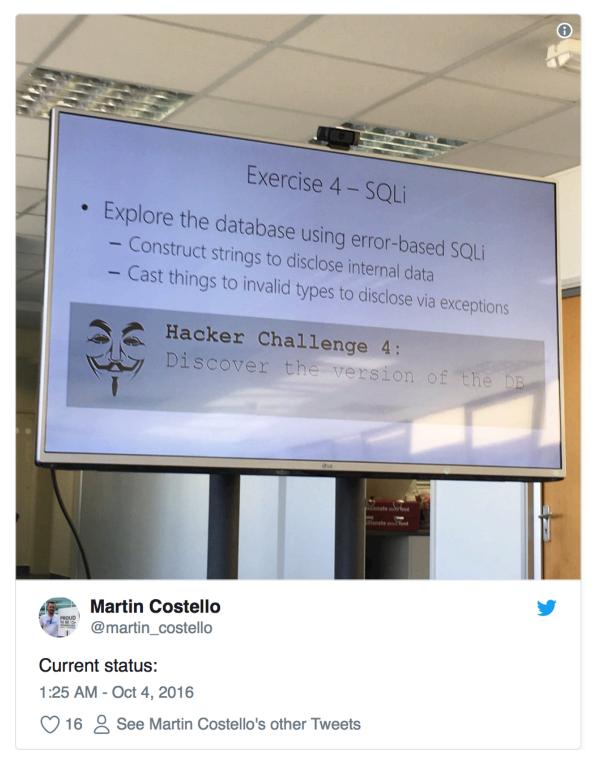
I'm up just after 6am for the first 2-day workshop of the trip (I'll do 5 in total).

It's a repeat customer who I'd previously seen at another location in Jan and they liked it enough to run it again in London for other team members (I actually had to bring the whole trip forward a couple of days for this).

It's another rare sunny day in London:



And from here on in, it's pretty much business as usual which means a non-stop 8 hours of talking and running through the workshop:

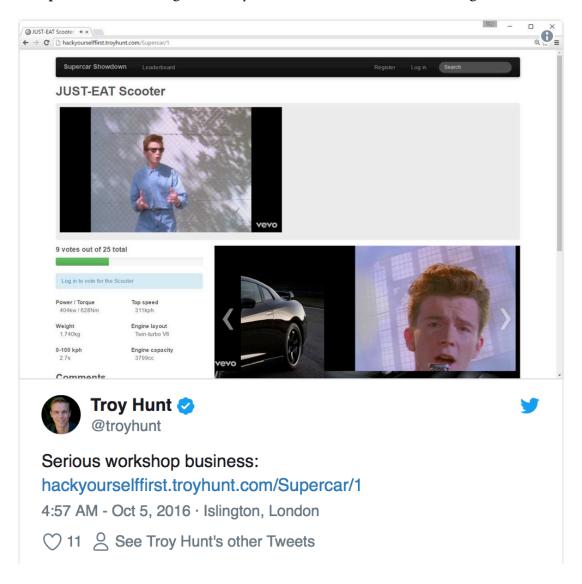


I go for some beers with the attendees later on, get room service back at the hotel and then crash about 20:30.

Day 4, Wednesday October 5: London workshop day 2 and flying to Edinburgh

I'm starting a half hour early as I need to battle London commuters to get to the airport for a 20:00 flight. Frankly, this is about the most stressful parts of these trips - rushing from one event to another where I *have* to make a flight or the next thing gets jeopardised. I've not slept great either - fine until about midnight but then tossed and turned for another 6 hours.

Workshop runs fine though, clearly those at the event are having fun:



As much as we have a bunch of fun in these workshops, I'm really glad to see people getting practical knowledge they can use in productive ways afterwards. I put a lot of effort into striking the right balance between engagement, entertainment and education.



We wrap up and it's tube then the Heathrow Airport Express train then security then lounge. It's about as easy as it can be (carry-on bags folks - that's the secret!), and I use train time to catch up on emails. Lounge time is to invoice the customer and again, work through email backlogs - there's no downtime. I'm extra conscious on these trips that if I start to fall behind, it's *really* hard to get back on top of things.

It's a short flight but it's 20:00 by the time I get on it and I'm fading. I don't want to sleep though as that starts to mess with the body clock and I'm still acclimatising. I watch some Breaking Bad on the iPad (re-watch it, that is) and stay awake. I get a tram from the airport then walk down unfamiliar dark streets on my own while towing my luggage and watching Google maps:

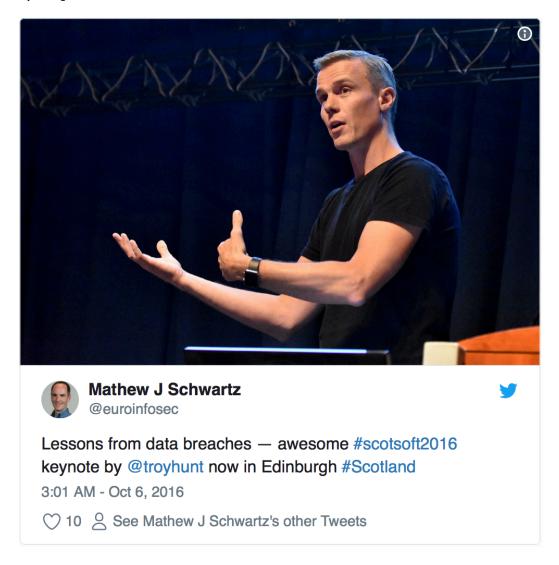


This phase of travel - the one where I've had a really long day then flown somewhere and tried to get myself to a hotel late in the day - is the most mentally taxing. It's just lonely and it's as far removed as possible from my family and home in the sun. By the time I get to the hotel and into my room it's 22:30. I'm seriously tired.

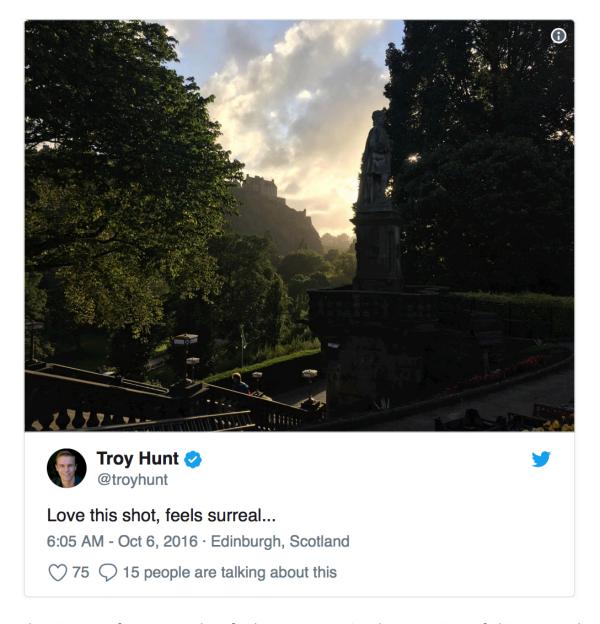
Day 5, Thursday October 6: Edinburgh talk then driving to Glasgow

I wake on my own at about 06:30 after an awesome sleep and feel really good. I need to rehearse the talk I'm going to deliver that day (*all* my talks get rehearsed multiple times), so breakfast is quick as is the daily family chat. By the time I do

all that plus check out plus get to the event it's 10:00 where I have a press commitment. And then another one. And then another one. And with a bit of socialising as well it's already time to set up for my talk, which ends up going perfectly to plan:



Per the earlier link, these only go to plan because I plan meticulously. I do some more press afterwards (seems to be a lot of these folks at the event) and then take a bit of a break. I try really hard to get out and about as well during these intense times; I have to get away from things sometimes and just step outside. And wow - outside was awesome:



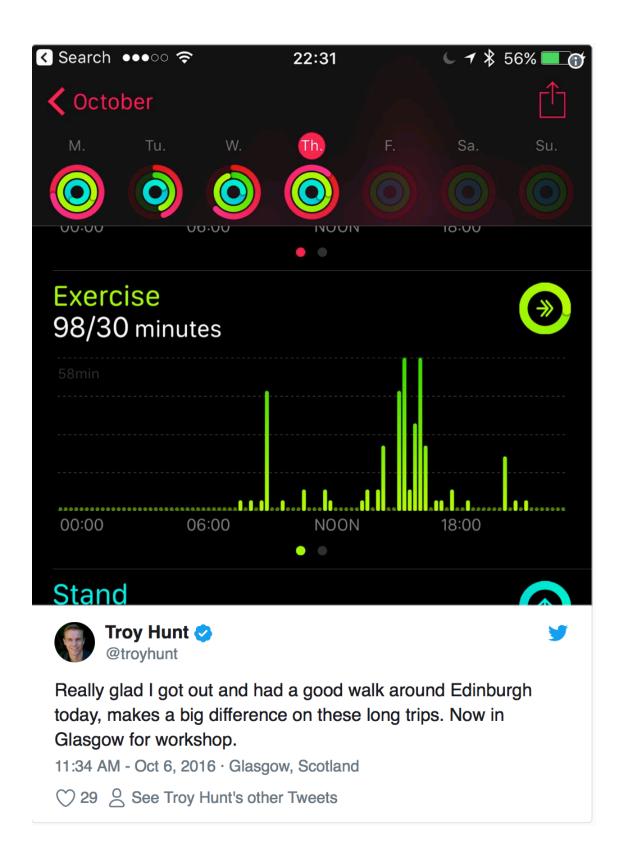
But this is a perfect example of what I meant in the opening of this post: that's an epically cool shot that looks great in a tweet, but it doesn't show how damn tired I was by then nor how absolutely non-stop the preceding three days had been. This was snapped during one short break before heading back to the conference. But before I did that, I also record my 3rd weekly update video, my first away from home:



I get just enough time to edit and upload the video after my walk then it's a conference dinner. There's always people that want to talk and believe it or not (and trust me, I'm still getting used to this), people that want selfies with me:



I keep it early because there's then a 1 hour drive to Glasgow (although fortunately done by my gracious host), check in to the 3rd hotel of the trip already and then a 22:30-ish bed. Long day but hey, at least I got some exercise!



Day 6, Friday October 7: Glasgow workshop day 1

Crap sleep. Tossing and turning from midnight to 05:30 when I eventually got up and I'm not sure why, but it means starting a bit behind the 8 ball today. At least being up early means time to catch up on a bunch of things. It's another private workshop today and I walk the 20 minutes to the office and get there by 8am.

This is a Friday / Monday workshop and I need to make an early departure on Monday to catch a flight so we've crammed more into this session. That means 08:30 to 17:30 and everyone's brains are pretty much mush by the end of it, mine included. So, we get beers:



It's great having social interaction like this; you have a lot of banter that wouldn't normally happen in the more formal office environment and it's fantastic for building lasting relationships. I'm also always conscious that it's very easy to have a few too many beers and really set back my sleep and overall health in a pretty intense work period so I'm out of there by about 9.

Day 7, Saturday October 8: Down day in

Speyside

The organiser of the workshop is taking me up to Speyside for the weekend so think picturesque Scottish scenery and whisky. Normally I'd be either flying somewhere or amusing myself on a weekend during a trip like this so it's quite unusual. We head off on a 3-hour drive and I start to see how Scotland is pretty cool;



Highways are boring though so I write <u>a blog post about Chinese data breaches</u> and get some more data loaded into HIBP. Without wanting to make it seem like I never tune out, when I come on a trip like this I'm considering it work from the

moment I leave home to the moment I get back so whilst I definitely have some downtime later on, I'm not passing up this opportunity to do something productive.

Speaking of downtime, we find a cosy spot later:



Also, I learn that trolling Scottish people is like shooting fish in a barrel:)

Day 8, Sunday October 9: Heading back to Glasgow

My son is having his 7th birthday party today. It's hard being away at times like this, particularly seeing video of my entire family celebrating whilst I'm on the other side of the world. That's the nature of travel though and it's very hard to work it around personal events like birthdays.

I go for a walk to console myself:



Then it's back to Glasgow again which means more hours of highway and some time to work away on the laptop. Connectivity is a bit dodgy on my tethered iPhone so I'm crunching through tens of millions of records locally within SQL Server running on the Lenovo P50. Every time someone tells you "oh, you don't need a powerful laptop, you just spin up a VM in the cloud", remember that there's nothing like being able to run queries on the metal using a Xeon processor and 64GB of RAM sitting on your lap without the need for an internet connection. I get a lot of stuff done on that trip:)

I get back to Glasgow around lunch, check back into the same hotel I left the day before then go for a wander. I've lost my only toothbrush and the one pair of jeans I brought have got a hole in them so both those are on the cards for replacement (the joys of travelling light). These are the sorts of things someone like my wife would never screw up but somewhere within my compartmentalised mind I've put them in the "not really critical" box where all the things I can easily fix on the fly go. I do get to see some nice sights though:



I get back and do a few hours of work, including parsing some more data breaches for HIBP and editing some Pluralsight. I've had a 14km walk today so I reckon it's ok to pop open a bottle of the home brew my host gave me (which was *sensational!*) and splurge on a burger:



I crash out early again because the coming week has the busiest schedule of the entire trip.

Day 9, Monday October 10: Glasgow workshop day 2 and flying to Copenhagen

It all starts to get a bit "Groundhog Day" here - same routine over and over again. Get up, hotel breakfast, walk to an office somewhere and do my security thing. It all goes totally fine with this one finishing a bit early so I'm off at 15:30 and being driven back to the airport in Edinburgh. I'm flying Ryanair this time and people have set the bar of expectation *very* low but right up until getting onto the plane, nothing has gone other than perfectly smoothly:





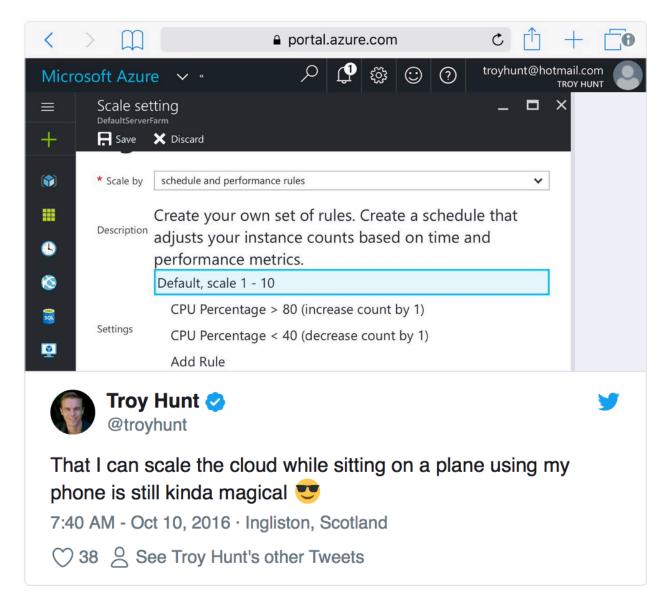


And that's it for Scotland this trip, thanks for a great time folks. Next stop: Copenhagen!

7:35 AM - Oct 10, 2016 · Edinburgh Airport

∑ 25 See Troy Hunt's other Tweets

And then I get to my seat... which is actually fine! Ok, it's not first class Qantas but it's right up the front on an aisle so even my long legs have plenty of room. I even wrap up a bit of cloud scaling after <u>loading the Chinese NetEase data into HIBP</u>:



I hit Copenhagen and get met at the airport by the workshop organiser and taken to the hotel. This is now the 6th time I've checked into a hotel already and it's only day 9. Fortunately, it is *awesome*:



There's a big lounge room, awesome bedroom, massive bathroom and big stone bath, then there's the rooftop; day bed, spa and a view over Copenhagen. I'm torn because on the one hand this is clearly awesome, yet on the other hand I'm there on my own without my wife, and there's going to be a *very* small number of hours where I'm actually awake in the place. Again, all this stuff can look epic in photos but the truth is frequently very different to the mental picture people frequently form.

Day 10, Tuesday October 11: Copenhagen workshop day 1 and .NET user group

It's the halfway point of the trip. It's also my son's birthday. And I'm not there. I

have a good chat to him on Facetime, but it's hard being so far away:



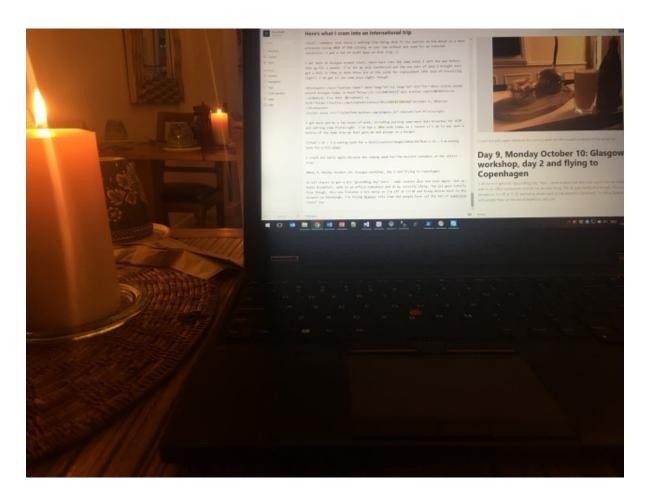
I spent most of my teenage years living overseas away from everyone but my immediate family. Dad was a pilot which meant not just the half decade on the other side of the world as a family, but him frequently being away on birthdays or at Xmas or other times families traditionally spend together. But you make it work in other ways and I suspect that's shaped my tolerance for not always being with my family when I'd like to be. Other people would never be happy

doing that and I totally get why.

I head up to my rooftop, first lamenting never getting to hop in the spa then snapping off a quick pic:



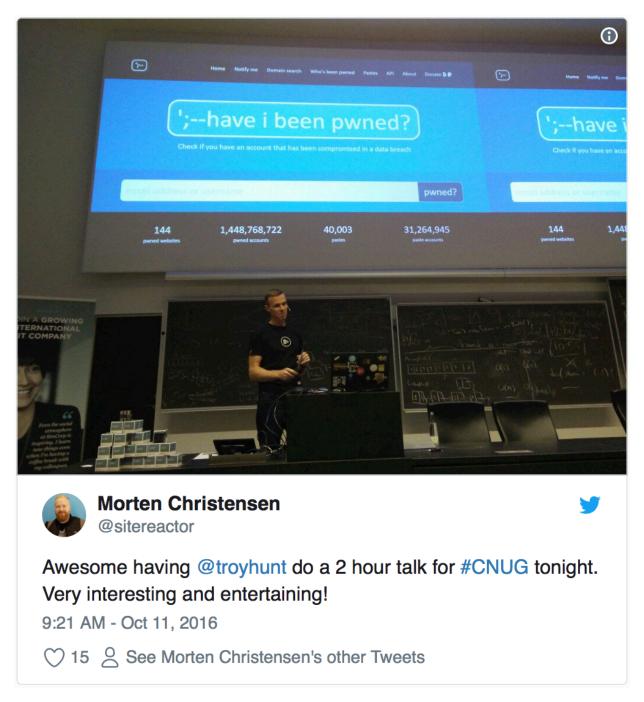
I'm at breakfast. It really is an awesome hotel; very Nordic yet somehow warm and cosy so I'm writing this by candlelight in the dark Danish morning:



More than anything though, I just appreciate having quiet time where I can do my own thing, even it's fleeting. I'm met at the hotel and then it's off to the workshop. Same deal as usual, same old spiel and fortunately, the same levels of enthusiasm from everyone. The day goes flawlessly and it's always great to see feedback like this:



I'm starting to feel run down though. Just a bit tired, a bit congested and feeling like I need some downtime. But that's not happening tonight, instead I'm doing a presentation to the local .NET user group:



The talk is the same one as I've just done in Edinburgh so at least preparation was simple. It's a user group in an auditorium within a university and it's a lot more casual than a formal talk you've got a limited time for so I embellish a bit.

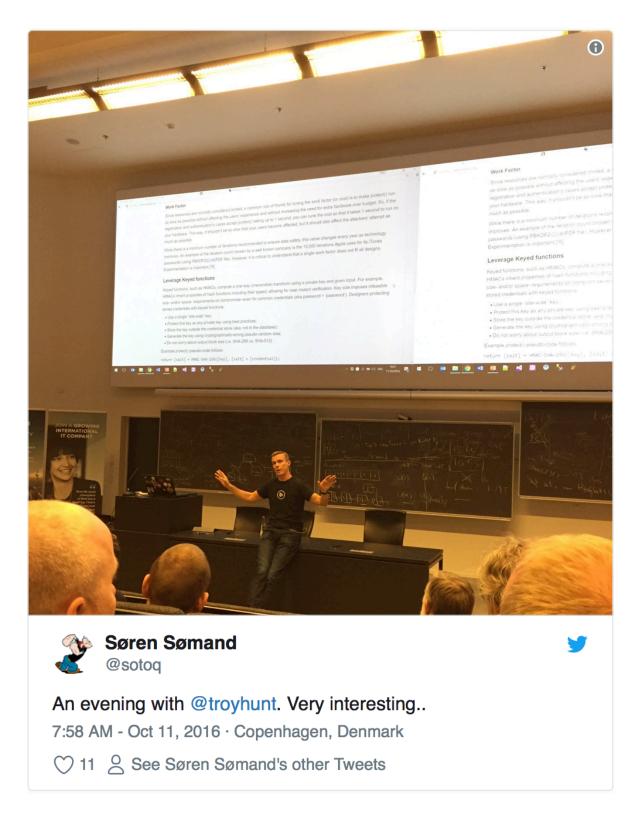
It goes for about an hour and then I spend another hour answering questions from the audience:



It seems Twitter hates images, here are some more from the @troyhunt talk.

10:36 PM - Oct 11, 2016

○ 13 ○ See Allan Kimmer Jensen's other Tweets



That goes great and I head back to the hotel with a few folks from the event when there's a suggestion of a local craft beer place. I'm tired, but I want to feel like I get to see at least a little bit of Copenhagen and spend some social time with people so two of us head out. The beer is rather sensational:)

Back at the hotel, I'm walking to the lift and the barman suggests I should *really* try some wine from a bottle he's just opened. I hesitate, then notice a sign which makes a lot of sense:



I don't know that I've ever just sat at a bar and talked to a bartender before, but

he was a lovely bloke and we talked a lot about travel. Zero cyber-talk or what I was doing there and that was just fine as it was nice to tune out for a bit. I still got to bed around 22:30 which was ok and I got a great sleep in my awesome room.

Day 11, Wednesday October 12: Copenhagen workshop day 2 and flying back to London

First things first - check this out:







This is awesome - the Danish user group I spoke at gave me a little LEGO me riding a jet ski chasing a cyber criminal :)

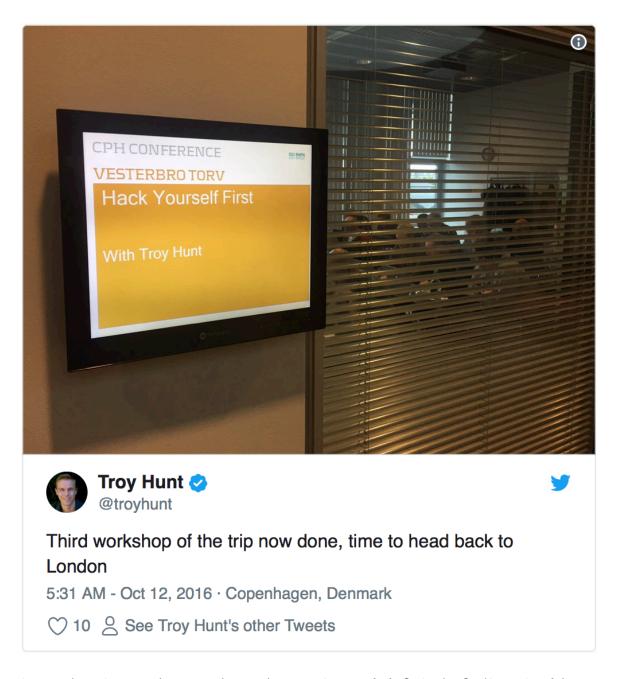
7:48 PM - Oct 11, 2016

○ 238 ○ 31 people are talking about this

This was such an awesome little gift, particularly from the home of LEGO! Little things like that really make your day so I'm happy. The hotel also gives me a bunch of organic shampoos and things in a goody bag as they're a very alternative sort of setup here. Unfortunately, large bottles of liquids don't mix with carry-on baggage so the workshop organiser's wife has done quite well out of my stay.

The workshop runs fine and I leave 30 happy participants:





Train to the airport, breeze through security and definitely feeling tired by now. Fortunately, Copenhagen has an *awesome* airport with lots of good healthy food which was just what I needed:





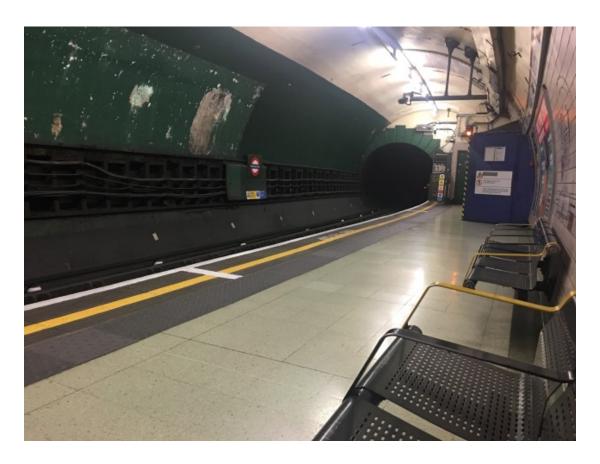


This is how to do airports folks, massive juice bar in Copenhagen and don't think I've seen a fast food joint yet

6:01 AM - Oct 12, 2016 · Copenhagen, Denmark

When I checked in earlier I upgraded my ticket for a small fee (seems quite cheap if you do it late when there's spare seats), so I got to relax with a heap of room and eat a meal once I was on the plane. But that was the end of luxurious experiences for a while...

I train it into London from the airport via the Heathrow Express which is pretty awesome then wait at a lonely tube stop in Paddington:



This is the part of travel that starts to get depressing; late nights, tired from the day and strange - or no - faces. I don't think this train has been cleaned since the great depression either:



By the time I get out of the tube it's after 22:00. Problem is, I can't find the hotel. It was on the map and I'm in the right location, but all I find is a door in a nondescript wall. But there's an intercom so I buzz and am let in. I walk up a narrow, steep staircase into a tiny reception area and realise that yes, I'm in the right place. The receptionist explains that my room is downstairs... in the cafe. She leads me down and back out the front door then along a few meters to another locked door which is indeed a cafe. We walk past all the tables and chairs, past the kitchen and through a "private" door. I'm in room "C" which I suspect they named after "cramped". Or "cooped up". Possibly "can't believe this is the room", who knows.

I definitely don't expect a standard like I'd just come from in Copenhagen, but I was pretty unimpressed. There was no desk to work from, no window (there's a blind with frosted glass behind it), no phone and definitely no room service and as I later discover, no iron for the shirt I need to wear the next night.

Day 12, Thursday October 13: London workshop day 1 and Pluralsight dinner

At least the room was quiet, other than the random buzzing that went off in the wall several times during the night, including at 05:30 which ended my sleep for good. I walk crookedly through the bathroom door, not due to the lack of sleep but because I'd hit my head on the low frame otherwise. Shower and grab an old fraying towel then have a quick chat to the family (they're beginning to feel less envious of my "glamorous" travel now).

Now here's the other problem: this workshop only has 6 people in it. There are various factors contributing to this which are all totally upstream of me and out of my control, but I'd normally have 5 times the number of people and unfortunately in this case, I'm paid based on profit share. Almost every other training event I do is a flat rate (those attached to conferences are the exception), which means that at least financially, attendance numbers have no impact on me. But this one will really bite and I'm as frustrated about the low turnout as I am about not using my time as efficiently as I could have.

But more importantly than that, 6 people have paid to come and see me talk for two days and above all else, they've gotta love this workshop. I still need to deliver the best possible experience to them and if there's one thing I never sacrifice on with work, it's quality.

Moving on, at least the coffee in the cafe-thoroughfare-to-my-bedroom is good and the workshop facilities are literally over the road. All 6 people show up too and despite the circumstances, I'm happy with how it's all gone. A couple of the attendees weren't very strong in terms of web programming ability (we do a bit of HTML and JS), so I pair people up which works *really* well. I must think about

this more for the future; I first paired people in the US a few months ago during a private workshop where there were 50 people and I wanted to keep the number of machines down so I could spend more time with people. For all the same reasons pair programming works well, pairing in the workshop makes a lot of sense; collective problem-solving, learning from each other and the exercises run a lot faster too.

We finished at 17:00 and I'm chairing a dinner put on by Pluralsight at 18:00 for some CISOs and other security bods they have relationships with in London. I brought one shirt with me for this event but I can't iron it because, well, let's not get into all the things this hotel room doesn't have again. I walk 20 minutes to the event which fortunately, is in a nice spot:



Seeing friends from the company is *really* nice after the way I'd been feeling since arriving in London. It's always awesome seeing Pluralsight folks across the world whether they be staff or authors like me; there's a great comradery and for

an event like this where I'm turning up to face enterprise customers, there's a sense of really being wanted. In fairness, workshops feel the same way too, but this is a welcome change from something that's rapidly becoming repetitive.

I finish up dinner and snap a pic of the hotel as I arrive "home", lit up in all its glory:



I enter through the locked cafe door hidden behind the stall in the foreground and enter the deserted cafe:



The door at the end of the hall leads to several rooms and I squeeze into mine, tired and a bit fed up. One week from now and I'll be on the plane home then shortly after, sitting in the sun in a place I love. I need to remember why I'm doing this, and having the life I have back home is a big part of it.

Day 13, Friday October 14: London workshop day 2

I wake to a message from my wife saying my son has hurt himself. There's a photo of his chin with a big split in it and blood everywhere. I call her but they're in the hospital and he's just about to get stitches so she can't talk. I check my email while I'm waiting for her and see a thread of problems relating to

maintenance issues with the house that could have been pretty serious. It's made things really hard on my wife and now she's got an injured kid as well. This is the part of being away that *really* sucks any remaining gloss off the notion that international travel is all glamour.

But these are all minor road bumps in the grander scheme of things and as much as they can distract you in the short term, you can't let them cause you to lose focus. Kids get hurt, stuff in the house breaks and so long as everything is fixable, let's move on and work on the things we can actually influence.

Back to day 2 of the workshop and it goes like clockwork. I do find myself embellishing a bit more as I do more of these and there's more news and other related stories to talk about. I'm spending a lot of time talking about Cloudflare in the HTTPS module I run not just because they do some very cool stuff in this space, but because their model raises many other interesting angles on the topic. For example, sites not protecting traffic back to their origin and being MitM'd (the Pirate Bay kerfuffle with Airtel in India is a great example), how we need to tackle the price and logistical barriers to going secure by default and who you should and should not trust to handle your traffic depending on your class of site (I question the logic of TPB using Cloudflare). I also talk a lot about defending against attacks by dynamically implementing firewall rules in Cloudflare when abuse is observed. If I'm honest, I'm a bit proud of myself with how well this model is working and people love seeing all the mechanics underneath the implementation. I'm happy to show and discuss things in a private setting like that I tend to keep out of the public eye too and it's very wellreceived, but does tend to eat into my schedule.

We wrap up and I head back to my hotel (just one more night...) then work on some HIBP features for a bit. Anywhere else and I'd be tempted to just order room service and chill out, but obviously, that can't happen here. Instead, I head out for a walk:



I grab some BBQ and wander around. It's pretty down near Tower Bridge and it was a good idea to get out regardless of the desire to escape the hotel. It also gave me a great spot to head back to the following morning; I haven't recorded my weekly update video and I really would like to keep them up. I've just gotta spend one more night in that hotel first...

Day 14, Saturday October 15: Leaving London and heading to Zurich

I'm up just after 06:00. I don't have to be, my flight isn't until midday, but the sooner I'm up then the sooner I'm out and I can focus on new things rather than lamenting the last couple of days. I head out with suitcase in tow and walk back to Tower Bridge (did I mention the importance of carry-on luggage already?). I'm there around sunrise and it looks *sensational*:







Nice out this morning

8:28 PM - Oct 14, 2016

Now I know I post a lot of awesome Aussie sunrises, sunsets, sunny beaches and so on and so forth and I admit, I do enjoy the reactions to them (particularly from my UK friends), but surely this buys me back some kudos with the Brits, right?! It really is gorgeous and it makes an awesome backdrop for my fourth weekly update:



It's also a good spot to Facetime with the family so they get to see a bit of it. I think back to my teenage years overseas just before the internet and with no way of communicating with family short of very expensive phone calls or faxes (yes, faxes). Every time I'm away and I can actually *look* at my kids I think back to that and remember how fortunate we are to have the tech now, even though one of the kids looks a little banged up:)

Off to Heathrow, and I love this:



So simple, but just enormously effective. I find the first class British Airways lounge (a perk of having a ridiculous amount of travel on Qantas and partners) which is a pretty bloody welcome change of style from the last few nights. The flight itself has me right down the back but for less than 2 hours, I really don't care, I just zone out and watch some TV on the Yoga 900 (it's an awesome screen for this sort of thing, much bigger than the iPad and folds open like a tablet so you can use it on take-off and landing).

Into Zurich, train it to the city and definitely find myself in the right place:



The hotel is close... and it's fine. It's certainly not Copenhagen levels of fine, but it's a proper hotel and whilst it's a small room, I've got a desk and all the usual

facilities you'd expect. I catch up on a few emails and other bits and pieces then head out for a walk. Zurich is nice:



I sit by the water for a bit and chill, but what I like even more than the visuals is the audio - there's a cacophony of *very* nice cars in this place! Ferraris, AMG Mercedes and a heap of high performance BMWs. Music to the ears!

I grab some dinner, but it's another lonely affair:



It's also expensive - spaghetti and a glass of wine is over \$50 Australian which is kinda nuts. I walk back to the hotel before 19:00 and see two separate incidents of drunk English tourists getting into fisticuffs within 60 seconds of each other, both spilling claret on the cobbled Swiss streets. Nothing like a cultured

European holiday...

I manage to rack up 15km of walking for the day which I'm happy with, but I should smash that tomorrow when *I have a whole day off all to myself*, the only one of the trip without having to be anywhere or travel to another location. Now *that*'s luxury!

Day 15, Sunday October 16: Epic Zurich walking tour

Pro tip: if you want to sleep in, don't leave your alarm on. So yeah, up at 06:30 but it'd be highly unusual for me to sleep longer anyway so I'm not too upset. I spend a few hours attending to the usual electronic things, most of it while sitting in the hotel restaurant enjoying coffee and breakfast.

I head out just after 10:00 with a fog hanging over Zurich. I've got nothing more than a vague idea of where to go, but I head off and snap pics of anything interesting along the way:







Oh nice, Switzerland has some animals that aren't poisonous, venomous, have massive teeth or drop out of trees on you

11:04 PM - Oct 15, 2016 · Zurich, Switzerland



It's going to be a big day distance wise, but it's also a day off and I want time to chill which means finding some nice sunny spots and just enjoying the place. I've got the Yoga 900 in my backpack too should I want to actually do something productive so I find a nice little spot to jot down some ideas:







Best coffee of the trip in one of the nicest locations

12:27 AM - Oct 16, 2016 · Zurich, Switzerland

○ 15 See Troy Hunt's other Tweets

And I did genuinely use that walking time to work on ideas too. I've had something in mind for HIBP for some time which I just haven't been able to properly position, but I reckon I've got it right now. I write it all down for later sanity checking and also manage to drag out an old blog post I've wanted to get out for a while and prepare this (I post Here's how I handle online abuse the following day). As much as I want to "tune out", I also like enjoying these quiet times to do the things I never seem to get around to while feeling rushed.

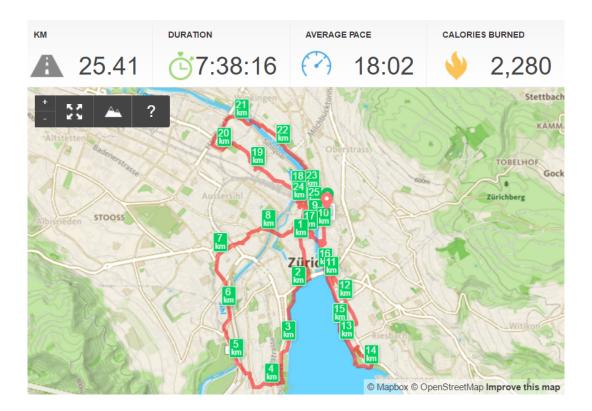
I eat in a nice little spot in the sun by the water then head back towards town to catch up with "someone from the internet". As with London, this is always a really nice way to meet local people and also as with London, it's a very positive experience and really adds to the trip. He takes me to a great local spot with a pretty epic view:



I head off for more walking, passing a pub with some rather tasty looking beers. I walk past... damn, they did look pretty tasty. *Really* tasty and I *have* had a lot of exercise today...



By the time I'm all done, I've covered over 25km which I reckon is a pretty good effort:



I've grabbed a salad along the way and a combination of lack of forethought and a relatively spartan hotel means eating it with my fingers. Such is my "exotic" jet-setting life.

Day 16, Monday October 17: Zurich workshop day 1

Back to work and the usual "breakfast with emails" routine. I'm met at the hotel at 08:00 and we wander out into the dark, misty Swiss morning. Train, setup at the office and then it's business as usual:





Christian Moser @moser_christian



Hack yourself first - before someone else does it. Security Workshop @zuehlke_group @troyhunt #cybersecurity

9:39 PM - Oct 16, 2016

30 \(\text{17 people are talking about this} \)

And it is usual - there's about 30 people in the room and we kick off the 5th and last workshop of the trip in tried and trusted fashion. Everything about the day went to routine which is just fine, but I find myself continually having to remember whether I've shown certain things already or made particular jokes or said other things which could easily have been said in a totally different workshop. Or this one. I'm not sure because I'm going through the same routine

over and over again.

That night - cheese:



A bunch of people take me out for dinner very close the hotel (big tip from me -

having events close to the hotel is *awesome* as I get to chill for a bit first), which is just a great night out. It's one of the most enjoyable evenings of the trip, just nice people and relaxed conversation.







Small but serious conversation about data leaks, kangaroos and sport cars with @troyhunt and @zuehlke_group

10:30 AM - Oct 17, 2016 · Zurich, Switzerland

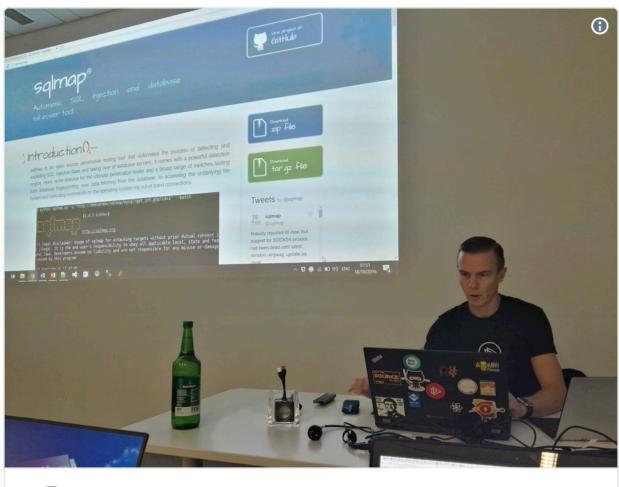
11 See Natalia Oskina's other Tweets

They also had a much healthier gender diversity than most places. Still far from where we'd all like the industry to be, but well above par and it does make a positive impact on the views and perspectives that are shared not just at dinner,

but throughout the workshop itself too.

Day 17, Tuesday October 18: Zurich workshop day 2 and train to Luzern

We're starting an hour earlier so I can get out earlier and head to the next event. I'm up at 05:45 and in the office setting up before 08:00:







Second day of the workshop and @troyhunt is ready for an early start

7:54 PM - Oct 17, 2016 · Schlieren, Schweiz

○ 16 See Daniel Rosendorf's other Tweets

Everything goes to plan, but by midday I'm tired. Really tired. I didn't have a late night or a lot to drink, but I'm conscious of how hard I've been pushing it for the last 2 and a half weeks. I don't feel so much run down or unwell, more that I'd just like to lay on my couch and watch movies for the afternoon. Clearly that's not going happen but I'm now pretty actively counting down the days and hours until I go home. In my mind, I've gotta get through until the end of tomorrow

night when I'll be in London again, talk there then I'm done. I've still gotta get to Luzern later in the day and speak at an event there tomorrow, but that doesn't help my mind trick of convincing myself that I'm almost done!

Regardless, I power through the day with some help from a couple of strong espressos (don't make a habit of this folks, it's not good for you in the long run), and get out a couple of minutes before planned close at 16:00. Brisk walk to the train station and the 16:04 back to the city is approaching. I *try* to buy a ticket from the machine but the train is here so... I figure I'll talk myself out of it if a conductor turns up: "Crikey, I needed a bloom'n ticket? How many dollarydoos is that?". But no conductor so a free 11-minute ride is all mine.

Back at Zurich station, I catch up with Scott Helme. You may remember Scott from such episodes as "let's hack his car and turn on the heater while he's freezing his arse off in England and I'm chilling by the pool":



I'm conscious by now that I don't even know if I'm going to Luzern or Lucern. Or is it Luzerne? I'm literally at the point where I'm just following TripIt instructions and not worrying about the details. As such, I'm happy just to follow someone who's a bit clearer of mind.

It's nice catching up with a friend and being able to cut though the small talk.

It's a one hour trip but it absolutely flies and we're off before we know it. We have a quick walk through what turns out to be a *very* picturesque little spot:

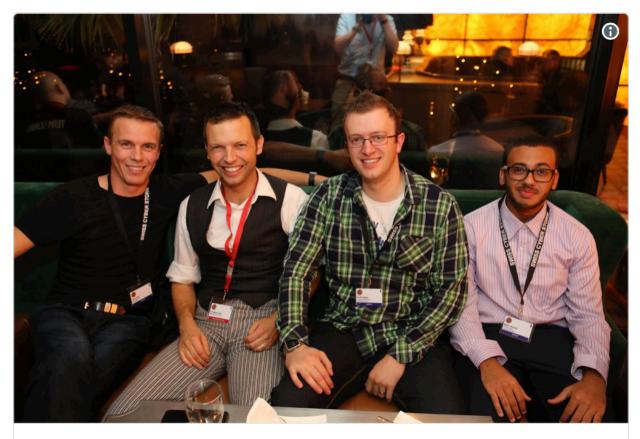


Check into the hotel, up to the room... and it stinks. It's *really* smoky and I honestly walk out and back in 3 times just to make sure I'm not imagining it and it isn't some weird Swiss air freshener or something. Nope, it genuinely stinks and while I'm here, what's with all the smoking in Europe? It's par for the course in a developing nation like China or Indonesia, but it always surprises me when

I'm back in an otherwise very developed part of the world and people are puffing away, *especially* when you're sitting outdoors in a nice cafe and it's wafting over from the table next to you.

Anyway, unhappy reception phone call, back downstairs, new room. I have literally 5 minutes (maybe 6) before I need to meet up with Scott and head out for the Swiss Cyber Storm conference dinner.

It's a casual dinner for the event with lots of *unfamiliar* faces. Many events I go to are frequented by the same folks, both speakers and delegates, but it's all new here. Still, there were a bunch of really nice folks there:



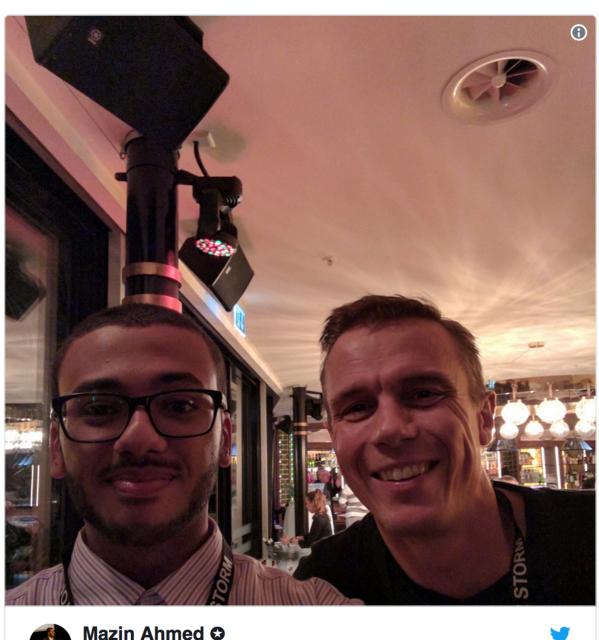




One happy @swisscyberstorm program chair / groupie with his speakers @troyhunt, @Scott_Helme and @mazen160.

3:28 AM - Oct 20, 2016

And as with other stops along the way, there was the selfie:



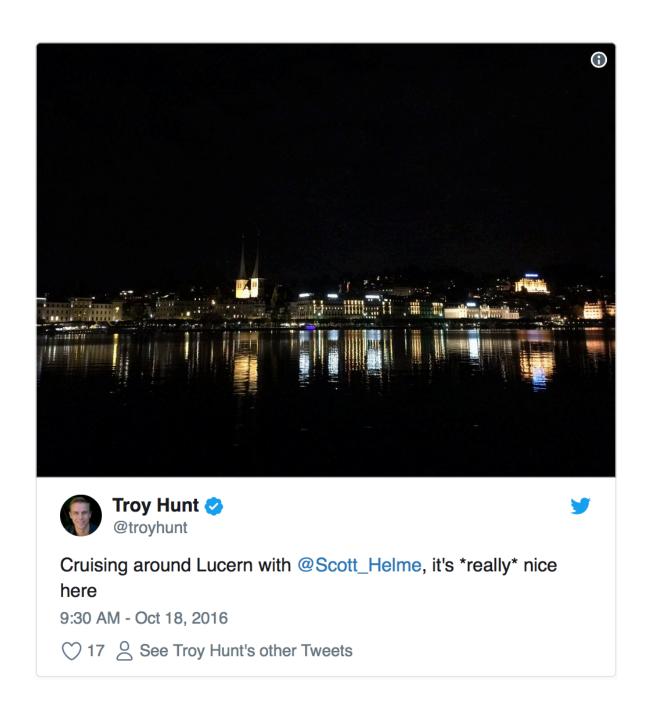




Had the chance tonight to meet @troyhunt in SwissCyberStorm! 12:22 PM - Oct 18, 2016

 \bigcirc 26 $\stackrel{\circ}{\simeq}$ See Mazin Ahmed \odot 's other Tweets

Wandering home, I get a bit of a sense of just how pretty it is here:



Day 18, Wednesday October 19: Swiss Cyber Storm talk and flying back to London (again)

I've had a crap sleep. *Really* crap. Couldn't fall asleep in the first place so ended up popping the first sleeping tablet I'd had since arriving then woke up at 03:00 and tossed and turned until 05:00. I get up anyway as I'm meeting Scott for a 06:00 breakfast so we can head out and see a bit of Luzern before the event.

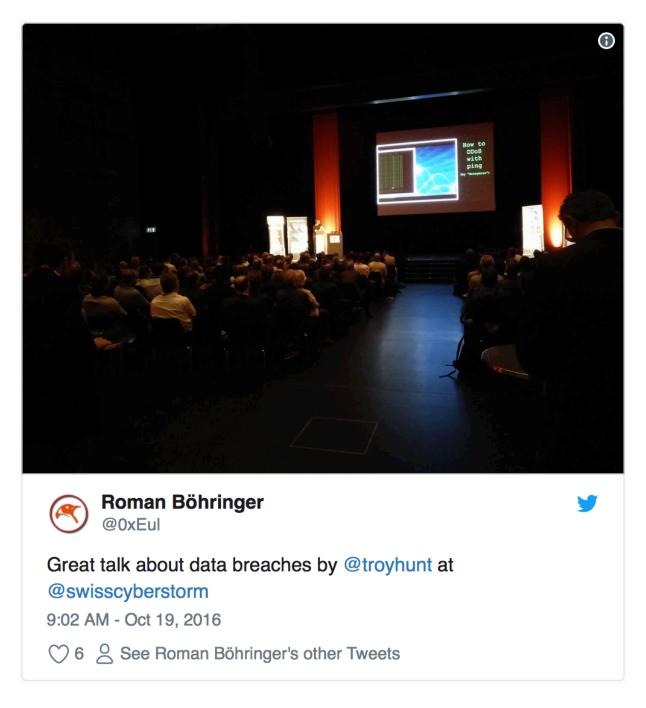
It's dark on our sightseeing tour, but it's also pretty awesome:



However, whilst walking around I talk to my wife. Son's stitches can't come out because he's having a reaction to the tape on his chin so the stitches have to stay in for a bit plus he's got conjunctivitis so now needs to be off school for the remainder of the week which means she's looking after him. It causes other complications as well, not least of which is her having to battle the traffic during the Gold Coast 600 which causes major delays around our house and she'll need to deal with it again tomorrow when the doc tries to take the little guy's stitches out again.

By the time I get to the conference centre where I'm speaking, it's finally getting light but I'm heading into a windowless auditorium. I get everything set up - on the stage, video perfect, audio perfect and the lot tested back to back over and over again. Everything is set to be perfect...

Then I get called on stage late. Then the video doesn't work. I'm standing in front of hundreds of people and only half my screen is visible. I'm mucking around with resolution and trying to make it all good whilst also trying to get the attention of the AV guys. There's nothing I can do but wait, so I'm making small talk with a darkened room of people that's more corporatey than my usual audience. Eventually though, it comes good and the show goes on:



It's quite a different audience to usual and I don't sense quite the same levels of engagement, but it might also be a more demure Swiss social norm. Regardless, from the feedback I do get (especially verbally afterwards), people are happy:



I get a chance to see <u>John Matherly</u> talk and for those who don't know him, John is the creator of <u>Shodan</u>, the search engine for the internet of things. Shodan often features in data breach and other related security stories as people discover all sorts of connected things that should never be there (MongoDBs with no authentication, for example). Scott and I catch up later with John and have a good chat; it's genuinely interesting work he's doing. The gears in Scott's head are obviously turning as he thinks of possibilities and this is one of the great

things about events like this: exposure to other really smart people during casual conversation that gets you thinking about things in ways you never have before.

I'm out mid-afternoon, pick up bag, train station, airport, hang around for a couple of hours (finally get some more Pluralsight edited), on the plane and into London City Airport (*so* much better than battling Heathrow):



I grab a cab for the short ride to the hotel (I'm totally not up for working out trains by now), and then - for reasons beyond all logic - have a taxi driver that

decides he can't take a credit card. FFS - you pick someone up from an international airport and then you can't take a payment with plastic?! It's such a minor thing in the grand scheme of life but the continued logistics of travel are really weighing on me. The hotel pays and bills me then I collapse into the tenth bed I've slept in since leaving home.

Then I realise that I *really* should run through this talk and get my timing down. It's a 15-minute version of the one I did in Edinburgh, Copenhagen and Lucerne which means a lot less slides and a very different pace. It's late, I'm tired, I've been going all day but I've just gotta do this one last thing...

Day 19, Thursday October 20: WIRED Security event and flying home

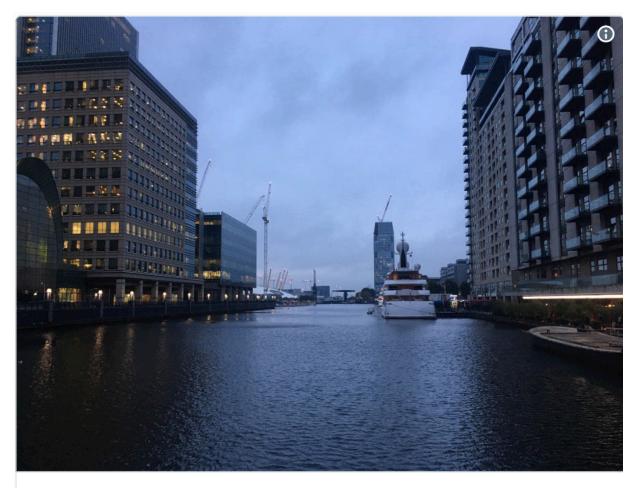
Up just after 5 but hey, it's the last day! I want to run through the talk again, catch the family, grab breakfast then get to the event by 07:45 where I'm having a breakfast catch-up with a bunch of people. Unfortunately, the laptop hasn't been on the charger (thank you dodgy international power supply adaptor), and I'm not sure when I'll get to add juice before tonight when I'll *really* need the power on the plane. But it's just one of those things that I need to put aside so that I can focus on the important issues of the day.

I get a drawing from my daughter while having breakfast:



It's been extremely hard being away, but I'm at a point now where I've made the mental switch to "I'm about to go home" so it doesn't pull at the heartstrings like it would have a few days ago.

I take a short 10-minute walk, leaving the last hotel of the journey and enjoying the last of the glorious British weather as I go:







Last day of the trip, wrapping up in sunny London

8:35 PM - Oct 19, 2016 · Canary Wharf Station

○ 19 See Troy Hunt's other Tweets

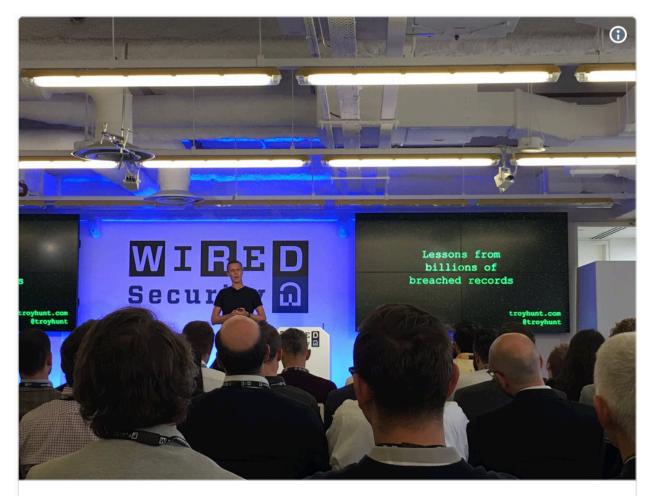
I arrive at the event which is *very* slickly organised. Very well-dressed people here too, making even my best t-shirt look a bit casual. It's the financial district of London so I guess it's to be expected and frankly, by this stage of the journey, it doesn't bother me in the least.

So, the event turns out to be awesome. Not just very well-run, but a cast of *really* top-notch speakers with genuinely interesting things to share. People like <u>Jamie</u> <u>Woodruff</u> who had a great talk on social engineering (not theory, stuff he'd

actually done), Moty Cristal who talked about negotiating ransoms demanded by adversaries who'd breached companies (the guy is an Israeli hostage negotiator - he's seen things!) and Mikko Hypponen who've I've spent a bit of time with in the past and really admire as a speaker.

Then there was <u>Mustafa Al-Bassam</u>, a quietly spoken bloke probably better known for the things he did as a member of the hacktivist group "LulzSec" a few years back than the positive things he's doing today. I had a chat with Mustafa after his talk and more than ever it struck me how so many smart kids find themselves at an infosec crossroad. He was only 16 when he and his cohorts were wreaking havoc and by any reasonable measure, deserved some serious repercussions as a result of their actions. Where the US in particular is throwing the book at people under the <u>CFAA</u> (<u>listen to Lauri Love</u> for a great example), Mustafa has faced penalties and moved on to become a smart, articulate and positive influence on infosec (he's presently undertaking a doctorate).

I do more interviews and then finally, there's my last formal commitment for the trip:





Zoë Rose #onHoliday @5683Monkey

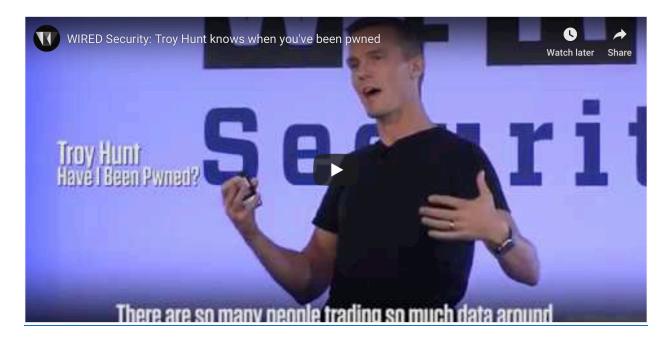


Stories from someone who deals with billions of records @troyhunt @WiredUK #WIREDSecurity

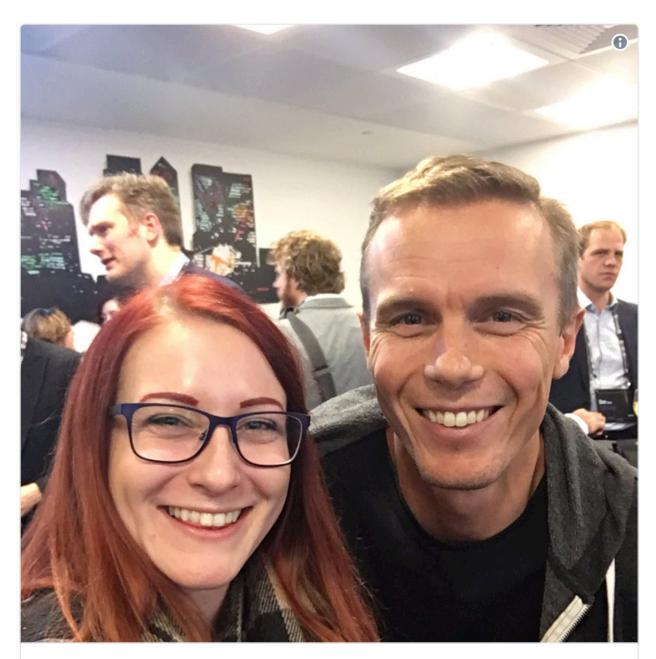
3:28 AM - Oct 20, 2016

○ 6 See Zoë Rose #onHoliday's other Tweets

It goes flawlessly and there's a massive relief from having now made it all the way through with no mishaps. WIRED runs a *really* slick event and they have <u>a</u> <u>full story on my talk published</u> within a couple of hours, including a snippet of video from the talk:



I spend time talking to a heap of different people afterwards which frankly, is the real value proposition of these events. People. Connections. Discussions which would never happen online, at least not in the natural, organic way they do at face to face events. Plus, selfies:





Zoë Rose #onHoliday @5683Monkey

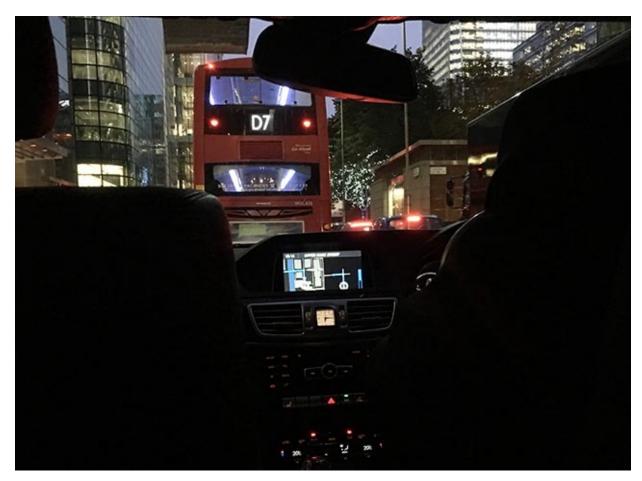


Replying to @5683Monkey

#selfiesForPer @troyhunt @thorsheim (Don't be fooled he had to basically sit on the floor.. #shortFTW) #WIREDSecurity 7:28 AM - Oct 20, 2016

○ 10 See Zoë Rose #onHoliday's other Tweets

This event has been absolutely spectacular; run with precision, fantastic talks, engaging conversation with delegates - I'm happy - this is the perfect ending to the trip. Come 18:15, it's time to play everything back in reverse from when I first arrived in London 2 and a half weeks ago, starting with the car:



The traffic is atrocious. It's going to take about an hour and a half to get to the airport but right now, I couldn't care less. I'm comfy and I'm not having to think. There's been a bomb in a tube station (or at least "a suspicious package") which has caused chaos that seems to be flowing over onto the road. The driver also seems to be having trouble seeing more than a couple of car lengths in front so he's on and off the gas like a crazy man.

Main thing is though that I'm at the airport with heaps of time to spare, fly through the check in and all the security bits and hit the lounge. There's massive relief to finally be done, but I'm ridiculously tired too. Take-off is about 22:30

and I'd normally be in bed by then at the best of times, let alone after the early start, long day and yeah, all the other stuff. I struggle to stay awake on the plane but I want to stay up long enough to eat. I sneak in one last weekly update video first though:



I kill a bit of time waiting for dinner, *attempting* to read some news via the inflight wifi. Thing about wifi on these Emirates flights though is that there's barely enough of it to even send a tweet:



By the time I get to sleep, it's 01:something back in London. I think. Who knows, any sense of time is about be thrown out the window anyway.

Day 20, Friday October 21: Landing in Dubai and leaving for Brisbane

I don't know how much sleep I got on the near 7-hour flight, but when I was woken up a couple of hours before landing for breakfast I was in total sleep-deprived zombie mode. I honestly can't remember ever having woken up this tired before. It takes me a good 5 minutes just to be able to sit up and focus my eyes let alone actually feel like eating anything.

I'm into Dubai just after 08:00. The wifi is frequently pretty awful in the airport plus per my tweet on the way over, there's no VPN allowed. I want to get the weekly update video I mention above loaded so I chance it relying "merely" on YouTube and Ghost's SSL (which is there for precisely such occasions when you don't trust the connection anyway...)

I've got a couple of hours in the airport before the next flight which means catching up on things again. The easterly journey home is always tedious because of the two short nights you endure; it breaks your sleep and I find it much harder to recover. I read a good piece recently which explains the science behind why this direction is worse and whilst they talk about the longer circadian rhythms being the root cause (and I'm sure that's a part of it), the more disrupted sleep patterns is what really gets me.

Regardless, I get a bunch more Pluralsight editing done once I'm back on the plane until I'm having trouble staying awake again. By this time, it's starting to get dark and in an attempt to acclimatise myself to the changing time zones, I try to sleep. I'll take a melatonin tablet about an hour before attempting to sleep for the next few days which is meant to be a more natural alternative to full on sleeping tablets and helps get you back into a normal sleeping cycle (at least in theory). Probably a combination of that and how massively tired I was to begin with helps and I get maybe 6 hours.

Day 21, Saturday October 22: Arriving in Brisbane and then finally home

I awake with a start and I kid you not, I was having a dream about security. In fact, it was about someone breaching my own physical security at home and stealing digital content which is pretty much a nightmare in my books. But oddly, the dream was clear enough that it focused on very specific things I've been meaning to do for some time and I reckon if I'm now having nightmares about them, I probably should get onto that.

I get back into Pluralsight and *almost* manage to finish editing the entire course before touching down. I'm literally editing the last clip of the course (there are 30 in total), so I'm pretty happy that I've managed to fill all available "down time" with something productive.

As when I left home, Qantas provides a pickup service so I'm out quick and in a waiting car on the way home. Particularly after being away for so long, fast-tracking it home is awesome. It also gives me a chance to have my scheduled monthly meeting with my Pluralsight editor. I'm conscious this seems like I might be overdoing it to jump straight off a tiring flight into a meeting, but it's a commitment I like to keep and the timing is perfect given this latest course is as good as done. While I'm talking, my wife tweets me:



I've said it before in <u>How I optimised my life to make my job redundant</u>, but the support of your partner is critical for this sort of thing. I've seen the stress it can

put on a relationship when your goals and expectations aren't aligned and it can be enormously destructive. I do what I do with the full support of a loving wife and we make decisions about travel like this together.

And then I'm home, far away from lonely airports and strange places, unfamiliar beds and cloudy skies. After seeing my family, the first thing I want to do is one of the simplest: sit down in the peace and quiet and have a good coffee:

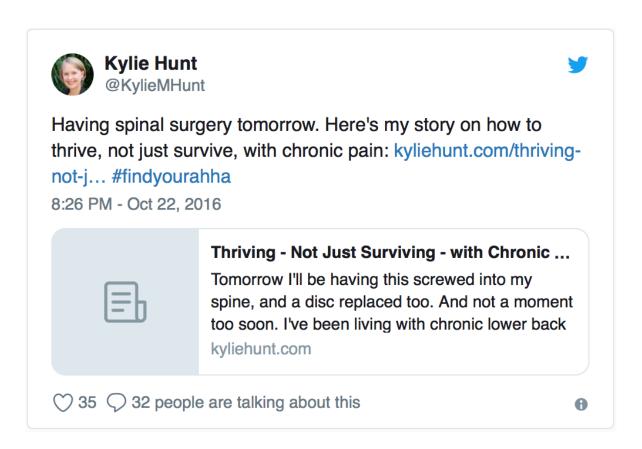


Post mortem

Let me point out something that I already knew, yet it became all the more apparent in reading back through this: notice how I've only tweeted photos of awesome things but amongst these are private moments that were sometimes really pretty unpleasant. We all do this - share the positive things on social media - but I really want people to understand just how tough going it was in amongst all this.

I've stayed in 10 hotels and taken 10 flights consuming 54 hours of flight time (that's not including all the waiting time and transiting between airports). I've missed my son's birthday, not been there when he's been injured and left my wife to deal with the lot for 3 weeks. And here's another thing that came as a surprise to many people when I wrote about online abuse last week: I didn't get paid a cent for any of the conferences or user group talks. Yes, people pay to attend conferences and no, tech speakers rarely receive anything. In fact, if you look at the comments there you'll see that some don't even get their expenses paid.

There's one other important thing I elected not to record as I went because I didn't know if she'd want it shared publicly, but given she's blogged and tweeted it I can share it here:



Kylie has had to deal with chronic back pain since Jan when coincidentally, I was also away in the UK. For the fifth long trip this year, she's had to balance the pain with not just all the household duties every family deals with, but do it solo and whilst in pain. We still made the decisions for me to travel when I did with this in mind and looking back we wouldn't have made the decisions we did any differently, but her condition certainly added to the emotional strain on both of us and *definitely* added to the physical strain on her as well. I'm finalising this blog post two days after arriving home and I'm very happy (and massively relieved) to say that the operation she had this morning went perfectly. I'm now looking after the kids for the rest of the week while she recovers in hospital so there's definitely no laying around to get over the jet lag (I've been up since 03:00 this morning).

Having said everything about how tough it was being away, on the flip side of it all I've done well from the commercial workshops and that affords me more choices now that I'm back home. That'll mean regularly taking the kids to school

and day-care, being there every time they play tennis and having lunch with Kylie each day. I suspect that many people would not willingly trade places with me given the family sacrifices involved, but this is a balance we have consciously decided is the right one for us.

I've now got a few months to do what I normally do from home which will mean more Pluralsight, working on HIBP, some local events and workshops and spending time writing about what I genuinely enjoy on this blog (I've a number of more technical pieces already in the works). Plus of course just generally looking after both Kylie during her recovery and myself whilst trying to get back into a more sustainable, healthy lifestyle, although I did actually manage to lose a kilogram whilst away despite the Swiss fondue (put it down to all that walking)!

I hope this has been an interesting read, I'm sure there will be those who both love and hate it for various reasons. If nothing else though, it's candid and honest and I hope it gives the reader some insight into what goes on behind the shiny travel tweets.

Comments

Hey - excellent post, and really insightful. I have a genuine question - you rightly point out that travelling is much easier when you've only got hand baggage... what, exactly, do you travel with? What sort of bags, how much do they weigh, do you count on using hotel laundry services and things? Do you actually fit your entire trip into airlines' baggage weight restrictions? I've flown with airlines before who restricted carry-on baggage to 7kg and found it a challenge to get everything in under the weight allowance - and if it's more than 2-3 days I'll just give up and check a bag. Genuinely interested in how you manage it.

Troy: I have a backpack with all the computer bits in the tweet towards the start

of this post and a wheelie bag that's the max allowable cabin size. They're choc full of everything which clothes wise is a heap of t-shirts and underwear rolled up into tight bundles. At worst, I wear some things twice but obviously you want to be a bit selective there! I can fit enough in that I don't need hotel laundry and frankly, particularly when the stays are often only one night plus the crazy price of laundry I don't want to be relying on that.

I don't think I've ever had my carry on bags weighed, maybe just once. Even then, allowances differ from airline to airline: I noticed the BA flights in Europe allow something like 23kg per bag even for carry on (which seems crazy) then Qantas says only 7kg. I'm always prepared to shuffle heavy stuff from backpack to wheelie bag though (stuff like the power brick for the P50) so worst case I can just check that in and possibly pay some cash, but that hasn't happened yet. Mind you, I don't try my luck with an airline like Ryanair or Jetstar here in Australia!

Wow, amazing trip! How do you plan all your flights and hotel stays? When I plan my trips (which are perhaps 1/10th as complex as this one) it seems like getting all the reservations and stuff is almost as arduous as the trip itself.

Troy: Very, very carefully! It all goes into TripIt and I run through everything meticulously in advance to make sure flights and events and rooms all line up. It's a lot of logistical work which I realise I haven't really captured here. So far, it's always gone like clockwork though.

"This phase of travel - the one where I've had a really long day then flown somewhere and tried to get myself to a hotel late in the day - is the most mentally taxing. It's just lonely and it's as far removed as possible from my family and home in the sun. By the time I get to the hotel and into my room

it's 22:30. I'm seriously tired."

I don't know if you looked into this already, but did you consider the Caledonian Sleeper for getting up to Scotland? Overnight train with beds, etc., so saves getting a late flight and a hotel like this. Sadly not as cheap as it used to be since the new franchise started, but still worth looking into for next time IMHO.

(Full disclosure: I'm a complete train nerd, so I would say this. I must confess I was wincing every time you took a car somewhere in the UK, but I'm sure it made sense in your case!)

(Also, that Bakerloo Line train you got from Paddington has probably had new seat covers since you went on it -- they've been refurbishing them recently, so it's understandable that it looks a bit worn!)

Troy: Frankly, that seems harder than it needs to be, the attraction of training it aside. At least flying it over quick and I still get most of a night's sleep in.

A genuinely interesting read. I rise to the bait of your perfect aussie beaches more often than I should;), but always in good humour. It's good to get a reminder, for me and others, that social media is only ever snapshots of someones life and they rarely post the true hardships. So many people compare their reality to someone else's highlight reel, without realising the comparison isn't only meaningless, it's unfair.

Troy: That's exactly it, although in fairness I guess that even before social media (or the internet, for that matter) we highlighted the memories we wished to keep as opposed to those we'd rather forget. It's easy to judge based purely on the glossy stuff though so hopefully this adds a bit of context.

Epilogue

Many aspects of this trip were routines I later repeated. For example, I had several trips back to Copenhagen where I stayed in the same hotel, although never had such an epic room again. I had a lot more London trips which were always a mix of having a great time with friends in interesting places, but also lamenting the weather, the tube, the weather and, yeah, the weather. I next went back to Switzerland a few years later and loved that, including ending up in Luzern again for a brief lunch whilst driving from Bern to Bellagio in Italy.

The family sacrifice is also something that resonates with me when looking back at the post, but not in any sort of negative way, quite the contrary. As I said in the post, the balance we struck at the time was the right one for my exwife and I and with the benefit of hindsight, I wouldn't change it. That balance gave both of us things we'd previously only dreamed of, not just in a monetary context, but in the way we lived our lives each day. For example, the ability for her to choose not to work is a luxury most people don't have, a luxury that was only afforded through the sacrifices described in the post. Same again when we later divorced; we're both able to live lives that are fundamentally different to if I'd been content going into a normal job 9-5, Monday to Friday.

COVID-19 fundamentally changed the way I worked and for a while there, put a complete stop to trips of any kind. If I'm honest, one of the things I fear most now that vaccines have rolled out and some degree of travel has resumed is that I'll get pressure to spend more time away. To some extent, I will, but I don't see me ever repeating the journey I described in this post. I dropped my kids off at school today, I'll see my family this evening and I'll walk on the beach with my fiancée on the weekend. I think I'm done with the epic journeys.

THOUGHTS ON THE LEAKEDSOURCE TAKE DOWN

They say imitation is the sincerest form of flattery, but **I hated LeakedSource**. Not because it could be viewed as a competitor to HIBP (why would I care when I was running it as a community service?), but because of the total disregard the service showed towards personal information. They couldn't care less about what happened to the individuals in data breaches and the more data they had on the more people, the more money they could turn it into.

I only touch on it briefly in this blog post, but I was also highly suspicious about at best, LeakedSource incentivising the data breach industry by paying hackers who'd breached systems and at worst, outright commissioning people to go and break into more systems. The flow of breaches during 2016 just didn't fit the usual patterns I'd seen over the previous couple of years and for some reason, huge amounts of new data kept turning up at LeakedSource before even a murmur anywhere else. It all just smelled super, super fishy.

I was also really worried about the negative light this service positioned HIBP in should it be perceived as operating in a similar space. I'd been really working on the transparency and ethics of HIBP to ensure this didn't happen, but it still kept me up at night. I deliberately avoided any public mentions or acknowledgement of LeakedSource to ensure I was never seen to even consider its existence, but behind closed doors I was watching every little thing they did. I was still watching it whilst at a conference in Belgium as the site turned from the familiar screen with a big search box for email addresses (wonder where that idea came from?) to the service being unavailable and news of the takedown hitting. I've gotta be honest - I felt particularly elated that day

27 JANUARY 2017

esterday, the website known as "LeakedSource" went offline. It's still early days and there's not yet an official word on exactly what happened, but the unfolding story seems to be as follows:

Yeah you heard it here first. Sorry for all you kids who don't have all your own Databases. Leakedsource is down forever and won't be coming back. Owner raided early this morning. Wasn't arrested, but all SSD's got taken, and Leakedsource servers got subpoena'd and placed under federal investigation. If somehow he recovers from this and launches LS again, then I'll be wrong. But I am not wrong. Also, this is not a troll thread.

LeakedSource provided sensitive personal information obtained from data breaches to anyone willing to pay for it. It was a service that occasionally popped up in news stories and <u>recently appeared in WIRED</u>. I've been asked for my views on the service in the past and how I felt about them providing passwords to people who didn't own them. If I'm honest, it's not really something I gave much thought too... until someone sent me my own personal data.

A few months back, a friend paid for their service and then went about notifying all their contacts who'd been put at risk. He sent me the following about my own personal data:

Filtered Results from DbForums.com
Hacked on: 2016-07-04
Encryption method for passwords: Vbulletin
username: troyhunt
hash: c7452c5009737e8b9cac
email: troyhunt@hotmail.com
register_date: 1359347497
last_login: 1359347805
birthday:
ipaddress:
salt: xgs`9t }5*"3pa3?]xzn:h'WoN]\[h

Filtered Results from VerticalScope Network (Vbulletin) (939 Websites)
Hacked on: 2016-02-01
Encryption method for passwords: Vbulletin
username: troyhunt
email: troyhunt@hotmail.com
ipaddress:
Real_Password: Reveal this result here
hash: 3e6fcf4bdce4913d8c119
salt: G1z
Website: Gtr.co.uk

Now let's be quite clear about this: I see a lot of data breaches in my travels but I was still shocked to see my own personal information sold in this way. My birth date. My IP address. My password hashes. My cracked password hash (a very old, very poor one). I know full well that my own personal information is out there in multiple data breaches, but there's a big leap between it circulating in relatively closed circles and being put up for sale for a few bucks. Not just that, but being sold on the clear web with no respect for the personal consequences of data breach victims such as myself.

LeakedSource appeared in late 2015 and quickly attracted controversy. Last year

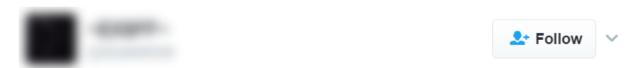
there was the cease and desist order from LinkedIn where they took issue with the redistribution of their members' passwords:

We have demanded that parties cease making stolen password data available and will evaluate potential legal action if they fail to comply. As a result we have sent a C&D to LeakedSource

A little while later, they were booted off Twitter with the account remaining suspended to this day. They returned under the guise of @BigSecurityNews but there was never any question that it was merely a facade to promote their activities in lieu of the primary account now being off air. Perhaps foreseeing the inevitable result this week, that account has been dormant for the last couple of weeks.

By late 2016, it was becoming apparent that their actions were erring very much on the black side of grey. There was a constant flow of data that wasn't appearing anywhere else in the usual trading circles before first coming to air via their service. Speculation was rife that there was incentivisation occurring not just to provide data that had already been obtained, but to actively seek out new targets that could subsequently be added to the feed of data then monetised by selling the personal information of the victims to whomever was willing to pay for it. This was always rumoured amongst those "in the scene", but it's not yet clear whether this contributed to the take down or if it was solely due to the services directly provided on the site.

There was never any doubt that the service was being used for destructive purposes. A quick trawl around Twitter shows just how other people's personal information was being used:



I wish leakedsource had an app so I can jack cod channels on the go!



i'm gonna do a ourmine used leakedsource and jack someone using linkedin's tweeting area



big thanks to @leakedsource & @dominos's for the pizza! yall both helped me get \$200 of pizza



what it feels like when you jack an account with leakedsource

```
| Grey Hat Hacker → Security Researcher →

XMPP: Security Is a Myth →

•Website Developer |
```

I've obfuscated these identities as I don't know what consequences those who paid for the service may now face, especially those who used it for malicious purposes. The theme here is very clear; the service was frequently used to do harm to others. Malicious use was broadly known and broadly discussed, even in the media:

The hacker will then run that username through LeakedSource.com and pay the website 76 cents for full results. In return, he'll receive an email address (which can be run through the database again for even more information) and password.

In fact, that piece went on to explain how other notorious hacking collectives regularly used the data to compromise victims' accounts:

Other hackers have stated that J5Z's LeakedSource method is the preferred strategy of OurMine, the collective that hacked Mark Zuckerberg's Twitter account.

And then there's the folks behind the service. Or the guy. Or girls. Or who really knows because from the outset, it was pretty clear they didn't want to be identified. The earlier mentioned WIRED article quoted them as saying:

if nobody knows who we are or where our site is located, bad people can't

attack us

The veil of anonymity provided them with a veneer of protection which was always a thin one. It's easy to see why they would have wanted to remain anonymous given the nature of the service and how everyone knew it was being used. However, even whilst operating from behind Cloudflare, the location of the service was readily discoverable by the casual observer and was always going to be easily accessed by law enforcement. There's a lesson in there for anyone who believes they can operate with impunity whilst trading on the misfortune of others.

Now, to Have I been pwned (HIBP). Some people have drawn parallels to services that both myself and others run:





So I wonder how the feds justified raiding @leakedsource vs @haveibeenpwned, @breachalarm, @lsLeaked, etc.

9:09 AM · Jan 26, 2017



I don't think that Mark genuinely considered HIBP or BreachAlarm to be operating in the same realm as what LeakedSource was (<a href="https://example.com/helare/he

HIBP never makes any sensitive personally identifiable data available to anyone, not even the legitimate owners of the data. In fact, some time back I wrote about how I will not provide data breaches to other parties either in full (I've never passed a breach to anyone else), or in part (I always point individuals to that post when they ask for their data). The only exceptions I can think of is when I'm verifying a breach and I've written publicly before about how I'll reach out to existing HIBP subscribers and seek their support in verification by providing them snippets of data. Certainly, under no circumstances would I ever provide

someone who doesn't own the data any access to it whatsoever. If you can demonstrate that you own the domain then you can see which accounts have appeared in which data breaches (many companies use this as a means of monitoring the risk their organisations are exposed to), but that's a far cry from handing over sensitive PII to strangers.

I've also never paid for data nor traded any of the breaches I've obtained. Creating a commercial market in no way improves the state of security, it merely provides incentive for malicious parties to obtain even more data.

Over the 3 and a bit years I've been running HIBP, I've found myself continually making small changes in direction in order to respond to changing sentiments and indeed changes in the data breach landscape. For example, when news of the Ashley Madison data breach hit, I elected to build out functionality to keep data from "sensitive" breaches beyond the reach of anyone who doesn't own the email address impacted by the incident (or the domain it sits on). At the time, that took a lot of thought but in retrospect the conclusion was simple: the data could cause serious harm to people so let's make sure that can't happen.

Some people expressed concern over them being discoverable in any data breach so <u>I introduced the opt out feature</u>. In that same post, I wrote about removing the VTech data because it was the right thing to do. Parents were able to get some comfort in knowing that data about their kids and indeed themselves had been removed from every possible location beyond VTech themselves and that was a very good outcome. In a similar vein, <u>I never loaded the Red Cross data</u> because like VTech, that was an incident we could contain and me not having the data was a very positive outcome (I didn't even retain my wife's or my own blood donation records from the leak).

I recently also <u>added a rate limit to the API</u>. I was seeing activity which I didn't believe was in keeping with the objectives of the project and posed the potential to put data breach victims at further risk so I put an immediate stop to it. (I later wrote about <u>how I blocked large volumes of malicious traffic using Azure functions and the Cloudflare API</u>, a model which continues to run beautifully to

this day.)

I'm writing this at the end of my second week abroad doing back to back talks and workshops focused entirely on trying to help people keep their data out of services like HIBP.





Just had a really nice two-day workshop on web security by @troyhunt! Oh btw, thanks for mentioning that Belgian bank KBC so many times ;-)



8:35 AM · Jan 26, 2017



I've said many times before that the best possible future for this service is that no more data flows into it; it would be an enormously pleasing result if more resilient systems were to stem the current flow of data taken out of vulnerable websites.

HIBP will continue to evolve. If public sentiment changes and, for example, the

premise of searching via email address become legally or socially unacceptable then I'll adapt. If other regulations require it, then I'll work to keep the service running responsibly and in a way that keeps both regulators and data breach victims happy. I think about this every single day.

As for whoever is behind LeakedSource, I hope this incident presents an opportunity to rethink the ethics of how personal data should be handled. The WIRED article stated that some of them "are still in school" and they may well just be kids who were attracted by the allure of some easy bucks without actually being malicious individuals. I genuinely hope the consequences aren't too severe and that they're afforded the opportunity to go on and do awesome things. For now though, the web just became a safer place by their absence.

Comments

Hi Troy. What if someone steals all the data you got on breaches. What would you do?

Troy: Well it would certainly be easy to load it into HIBP! I'd have to obviously disclose the incident and I'd notify my subscribers which would mean sending a million emails out. It would be impossible to notify everyone in the breaches themselves, I just couldn't feasibly send ~ 1.5 B emails.

The actual impact itself would be minimal though. There are only email addresses in the system so no other personal data to lose. That was one of the things that always worried me about LeakedSource; a breach of the service would have been very damaging for those involved. I don't *want* to have an incident on HIBP, but I've consciously taken steps to ensure that if I do, the scope of damage is limited to email addresses.

--

"There are only email addresses in the system"

Well, there's also a listing of which breaches are associated with the email address, which a

malicious actor could use to narrow down where to look for passwords that are reused. But that does require finding the other breaches themselves. Just considering the whole threat vector.

Troy: Yes, the association with the incidents would also be there and that poses some risk. The broader point I'm making here though is that I've made a conscious decision not to store any additional data attributes such that the worst case situation is the loss of the addresses and as you say, the association to the incidents.

--

It is a bit funny, at least for me that is, about all these breaches. If I keep featuring in loads and loads of breaches, then I just become immune to it to the point that I will start ignoring it. But if I do not feature in many breaches but maybe one a year or one in two years, then I will give it more attention. It's funny how it works. So for HIBP it would be a good thing if these breaches dwindled down, otherwise, people would just say "yeah, ok, there is another one, i don't care anymore. i will just keep using strong and unique passwords and forget about it". ...just what it is going round in my head...

Troy: I do wonder if we're suffering from "breach fatigue" sometimes, that is we're becoming somewhat desensitised and new incidents simply aren't having the impact they once did.

--

Leakedsource actually was BETTER then haveIbeenpwned in a few ways. Alerted me to breaches MONTHS before they were added to the pwned database, such as the Flashflashrevolution and Nihonomaru breaches, among others that STILL aren't in that database. I will miss it.

Troy: As this incident continues to unfold, earlier suspicions about their role in actively incentivising individuals to not only circulate data but attack sites to obtain it are coming to fruition. So yes, they would have notified people about breaches very quickly if they had an active role in the attacks!

True, but a lot of the leaks were months old, some years old. I would not have known about the flashflashrevolution or Nihonomaru breaches so early without it. I was able to alert flashflashrevolution back when that happened, but they failed to take attack until YOUR site confirmed it. So who knows how long the data was out there.

Troy: Often, data would be offered to me at a price and I'd always refuse due to the ethics of it. When a market is created for a commodity like this, people are incentivised to provide a supply and that makes the whole state of security worse for everyone. Leaked Source paying for data in this way always put them on thin ice and now they've met their inevitable fate.

Whether it's HIBP or any other service, we've only ever got a small snapshot of the larger breach landscape. I'll try and bring that to people's awareness as early as possible, but it's going to be done in a responsible fashion.

Epilogue

It was another year after writing this before I saw Jordan Bloom's name in the news. He was 27 by that time which I imagine puts him at around 24 when he got serious about trafficking identity information, a charge he was hit with alongside "mischief to data", "unauthorised use of a computer" and "possession of property obtained by crime". That last one worried me a bit as it's the entire basis HIBP is built on. But then again, most major tech companies had a bunch of the same data, as did identify theft companies and a raft of others using it in positive ways.

One thing I found particularly amusing about Jordan's arrest was that he was obviously trying to fly under the radar and not be personally tied to the site. Nowhere on LeakedSource could you find any information whatsoever about who was behind it, and he'd obviously made a very conscious effort to remain anonymous. Yet there in the news, a young Jordan Bloom is standing in front of A BRIGHT GREEN LAMBORGHINI! FFS Jordan, if you don't want to stand out

and look super suspicious go and buy a Prius or something!

LeakedSource wasn't the last service to suffer this fate, We Leak Info followed a few years later. Had they not read the news? Did they really think that standing up precisely the same sort of service would somehow lead to a different end? I suspect the problem here is a combination of huge volumes of freely available breach data combined with dirt cheap cloud hosting combined with young men seeing dollar signs and not yet having the maturity to think through the consequences of their actions. (Ok, that made me feel super old, but I think there's a great deal of truth in that statement.)

A couple of years after this post, Jordan plead guilty. I still don't know what his sentencing entailed, and I couldn't find anything when searching it whilst writing this epilogue, but I'm guessing it's not been a fun time for him. I'm also guessing he won't be the last to try this sort of thing on and just like We Leak Info after them... and literally while writing this piece, a reporter sent me a link to another site doing exactly the same thing. It even has "Leak" in the name of it! I'll get the popcorn.

RECKON YOU'VE SEEN SOME STUPID SECURITY THINGS? HERE, HOLD MY BEER...

This just may be one of my favourite blog posts ever. Throughout my writing and presenting, I've always tried to make what I'm communicating about much more than just the content alone. There's no lack of awesome content out there which covers a huge amount of the things I've subsequently written about, but what I'd like to think sets my writing apart from the rest is the show. It's the art with which it's presented, the way it's communicated and the emotion it leaves people with when they watch me talk or read my content. Humour and laughter, to my mind, are the best tools I have. I love it when I see an audience laughing and enjoying the show, I love that even more than the actual content and the knowledge they walk away with. When they go home later on, that's what they remember - how the content made them feel.

This post gave me the perfect opportunity to inject humour into topics that could otherwise be very dry. I just picked the stupidest shit I could find and dumped it into one killer blog post. The fact it got a couple of hundred comments is testament to how much people loved reading this and it gives me great joy even now, years later to read it again \bigcirc

28 APRIL 2017

y mate Lars Klint shared this tweet the other day:

spacemoses1337		
hunter		
hunter	u/rowealdo37189. Please choose a	nother password
email		
✓ I understand the risks of us	ing someone else's password	
I'm not a robot	reCAPTCHA Privacy - Terms	

@larsklint

Your password is not unique.

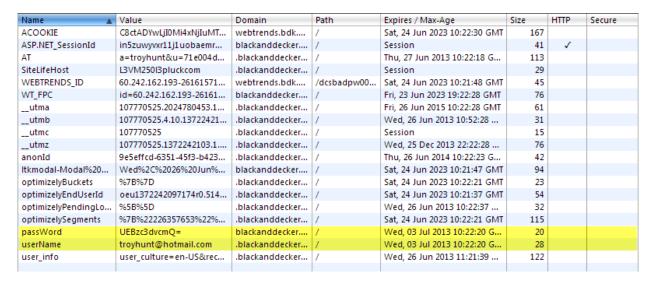
○ 2,669 9:17 PM - Apr 15, 2017

Naturally, <u>I passed it on</u> because let's face it, that's some crazy shit going on right there. To which the Twitters responded with equal parts abject horror and berating comments for not having already identified this as <u>a joke circulating on Reddit</u>. But here's the thing - it's feasible. No really, I've seen some *very* stupid security stuff out there the likes of which make the above example not just believable, but likely. Don't believe me? Here, hold my beer...

Remember me

Let's say you want to build a "remember me" feature, you know, the one where

you tick the box and then when you come back to the website next time, you're already logged in. Here's how Black and Decker did it:



Yes, that's just a Base64 encoded version of your password in a cookie and yes, it's being sent insecurely on every request and also yes, it's not flagged as "secure" therefore it's being sent in the clear.

Reckon that's bad? Try Aussie Farmers direct who I mention in that same post:



Oh wow, it's secure! But it's still a password in a cookie and it's still not HTTP only *and* they had reflected XSS risks on the site. And how did they respond once advised? That brings me to the next point...

Corporate responses

I did the dutiful thing and let Aussie Farmers know about the risk all the way back in 2013. I also suggested that maybe they shouldn't be emailing passwords

around (amongst a raft of other very nasty things) to which I received the following explanation from someone with "Marketing Manager" in their title:

To date we've not had a single security issue stemming from new customers being emailed their password, and I know for a fact 90% of the sites I personally sign up to online also follow that same process.

That reminds me of <u>this comment by Oil and Gas International</u> I referenced just the other day in the post on my new HTTPS course. This is where they got cranky because Firefox is now warning users when a login form is loaded insecurely:

Your notice of insecure password and/or log-in automatically appearing on the log-in for my website, Oil and Gas International is not wanted and was put there without our permission. Please remove it immediately. We have our own security system and it has never been breached in more than 15 years. Your notice is causing concern by our subscribers and is detrimental to our business.

Their website kinda, uh, stopped working not long after that (the SQL injection risks probably didn't help). They're back now, although it's unclear whether or not they've reset the clock on the whole "15 years" thing.

While we're talking about nonsensical security comments, <u>British Gas have</u> <u>struggled a bit in the past too</u>:



Pascal » Queue Jumper « Hartig @passy · May 5, 2014



Replying to @BritishGasHelp

.@BritishGasHelp Disallowing pasting and therefore password managers is NOT a standard practice. It's unnecessary and dangerous.



British Gas Help 🤣

@BritishGasHelp

@passy We'd lose our security certificate if we allowed pasting. It could leave us open to a "brute force" attack. Thanks ^Steve

○ 167 11:59 PM - May 5, 2014

0

And while we're in that corner of the world, it's hard to look past Tesco for <u>an</u> <u>example of corporate lunacy on the Twitters</u>:



But hey, secure password resets are hard! No really, check this out...

Password reset

It all started with this:

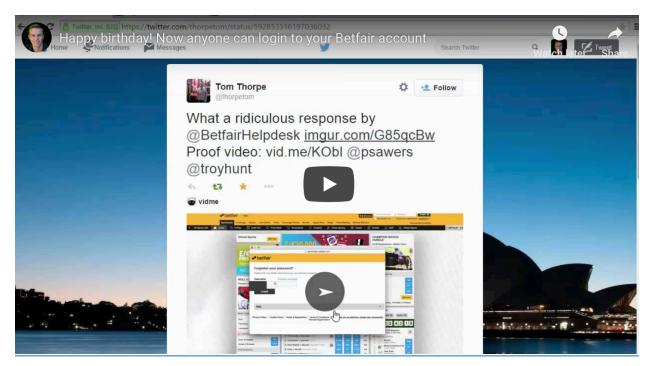




@BetfairHelpdesk Is it right that all one needs to change their password is their username and date of birth?

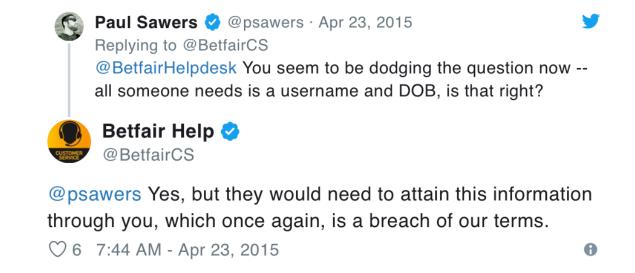
0

Now you may be thinking - "Oh, your username is your email address and Betfair would just send you an email and you'd reset the password via a unique link" - but by now you know that thinking would be too logical to make an appearance here. But amazingly, Betfair didn't actually believe Paul, so I made a video explaining it:



And it was exactly what it sounds like - if you knew someone's email address and birth date, you could reset their password to whatever you'd like it to be. But the pièce de résistance came with this exchange where, with what I assume

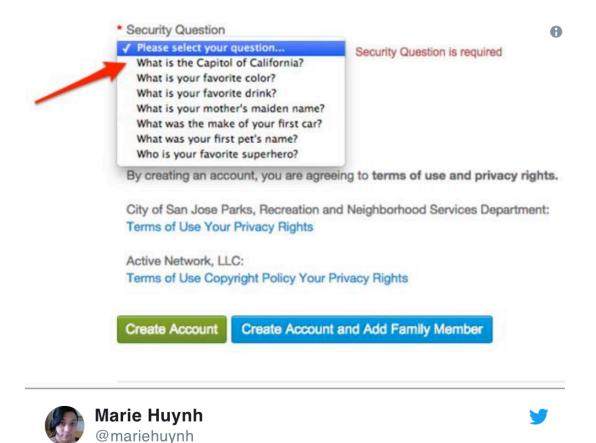
was a straight face, Betfair kindly advised that Paul would be breaching their terms if he gave his email address and birth date to anyone else:



You know what they really need here? Security questions...

Security questions

I'll just leave this one right here:



This is one of the worst security questions I've seen.

What? Too general? Try this one instead:



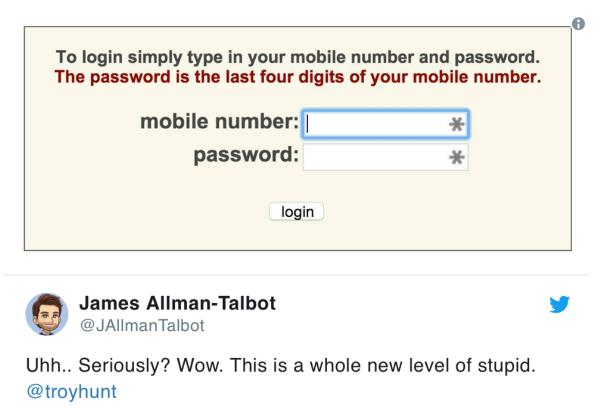
A security question on a website I was just on: "What is the name of your grandmother's dog?"

The fuck.

Because security questions are nuts! I mean those ones are extra nuts but in general the whole idea of taking either immutable pieces of data like your mother's maiden name or enumerable questions like the make of your first car or transient ones like your favourite movie... just the idea of security questions deserves a place in this post! Let's try something more sane...

Logon

You know what's hard? Passwords. If only there was an easier way:



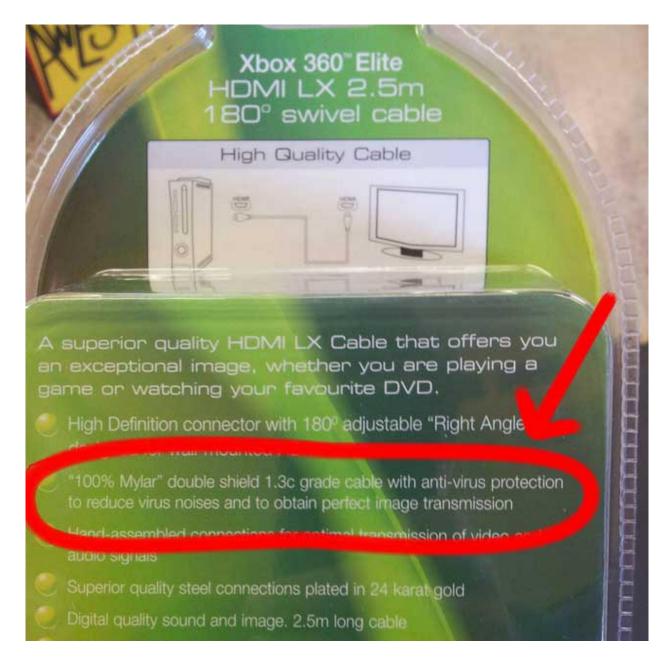
And before you go "but this is just a tweet and it may not even be real", it was real and here's the archive.org snapshot of it:

	Welcome to your SMS Diary - your own personal record of the texts you've sent us during the survey.
MY SMS. DIARY	To login simply type in your mobile number and password. The password is the last four digits of your mobile number. Login mobile number:
	password:

And before we all lose out minds going "the password must die", nobody has yet figured out how to make that happen! There are lots of *technical* solutions that nobody actually wants to use, the simple fact is we've got more passwords then ever and they're not going anywhere. But hey, I've seen worse...

Physical security

There's not really a way to position this without it seeming any more absurd than it already is, so let me just throw it out there:



You know the thing that really gets me here? Think about your non-techie friend and relatives who are just trying to get the TV and the DVD player working together. They go into the shop, pick up two HDMI cables and flip to the back of the boxes. They're comparing the specs - one of them has anti-virus protection and the other doesn't - what are they gonna do?!

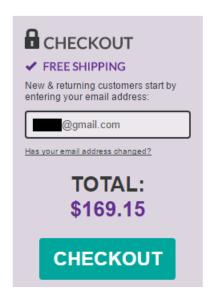
Now, just one more thing...

Account enumeration

I wanted to save the best until last. It's the best because it's still an active stupid security thing and it's inconceivably stupid but hey, at least they're fixing it:



Except that as of the time of writing, that was 8 months ago. And what is this stupid security thing? Well imagine this: you go to <u>Strawberrynet</u> and chuck some tonifying lotion or dry teasing dust or other thing I have little concept of into your cart then hit the checkout button. You're now presented with this:



So you enter an email address - any email address with an account on the site - after which you're presented with, well, someone else's personal data:

≜ EXPRESS CHECKOUT



Wait - what?! It's exactly what it looks like in that they'll hand over the personal data of anyone with an email address on the system. There's plenty of people on there too because they're within the top 5k largest websites in the world so you can head on over, enter a female name (they're largely selling cosmetics) then a popular email service and there you are! And in case you're thinking "well this is just terrible", no, it's actually a feature:

Please be advised that in surveys we have completed, a huge majority of customers like our system with no password. Using your e-mail address as your password is sufficient security.

No it's not! And no they don't! <u>I wrote about website enumeration insanity</u> back in August which is what promoted their earlier tweet and they appear to be completely oblivious to the problem. I even created an account myself just to check how it works:



I think I need another beer...

Comments

I just picked 5 email addresses from my inbox that belong to some female friends. 3 out of 5 of them had accounts on Strawberrynet and so it showed me their full names, phone numbers and addresses. How can Strawberrynet not realise how bad this is??

Troy: Please be advised that in surveys they have completed, a huge majority of customers like their system with no password. Using your e-mail address as your password is sufficient security.

I look forward to seeing the response by the UK ICO or another EU regulatory body if this isn't fixed by May next year ... €20 million euros potential fine under GDPR will no doubt get some attention.

Troy: Yes, I was thinking the same thing myself, they definitely do business in the EU too so extraterritoriality laws should kick in there.

I had a good laugh, but about the "anti-virus" cables: obviously this is a translation error. Cable shielding reduces "parasitic noise", not virus noise, but the person who did the translation probably missed that.

Once I saw a Lightning cable in a store being advertised as a "light up" cable. People doing the translation are usually not engineers.

But..... HDMI is a digital spec, if there is noise in the bitstream, the handshake likely won't

complete. It's very much a - you get a picture and sound or you get no signal at all - there is no middle ground with HDMI. So if someone buys a HDMI cable, and it doesn't work - it will be returned to the shop.

It's like the ridiculously stupid things you read about oxygen free cables (for domestic use) - that allow the electrons to move faster, or even stop the electrons flowing in the wrong direction. These can't be translation errors.....

Of course, it's either all or nothing, but I've seen cables that give weird problems. For instance, my provider's set top box uses HDCP to protect the content, but the overlays (like subtitles) and audio are still there. That took a while to figure out it was a bad cable.

... no, it is definitely a thing. The video data is not sent on the same line as display data. Noise can also be intermittent, so your handshake could even complete but your video is still noisy.

Either way, noise on any communication line is undesirable.

God I'm sick of this idiotic misconception.

Uhm, no, I think the overlay is superimposed by the set top box and in the main video stream sent out.

I thought this too, until I've seen with my own eyes a bad HDMI cable that actually added static to the video image.

Its the virus man... I'm telling ya, you need virus protection cables!

Hi Troy, you might get a kick out of this one too.

The PC Plus rewards program in Canada had as password requirements: "at least one uppercase letter, one lowercase letter, one number, and between 6-8 characters." As in, passwords shorter than 6 characters and longer than 8 characters were forbidden.

Shortly after their database got hacked, they changed the password requirements to this: "at least one uppercase letter, one lowercase letter, one number, and exactly 8 characters."

I asked them what the logic was with this. They responded by reporting my Twitter account.

I contacted them on their contact form after there was a request to reset ALL account passwords following reward points theft for many users:

Reason:

Technical Inquiry

Subject: Password requirements gone wrong

Message: Hi! Your programmers and IT staff should read this and repair

the password mess that not following the rules in the document created. $\,$

Link:

https://nakedsecurity.sopho...

FYI, all my passwords are unique. I understand and know not all users follow those practices, and I congratulate you for warning your users about good password hygiene. But your main failure point for me is limiting password length to 8 characters. Go read that link, pass it to IT, Have it explained to Management. You have a chance to get ahead of the crowd, just expand the changes you are currently doing. Oh, and please, don't be another Yahoo! of security!;-)

Martin Boissonneault

Working with Computers since 1990.

Their reply:

Dear Martin,

Thank you for contacting PC Plus Member Services.

We're always looking for new ways to better suit the needs of our customers. That's why feedback such as yours is always encouraged and appreciated.

The updated requirements for passwords were carefully considered and took existing PC Plus member feedback into account along with increasing the existing requirements for password complexity. While PC Plus passwords remain 8 characters, they are now required to contain an upper-case letter, a lower case letter, and a number.

Please note that special characters (i.e. !, ?, #, \$, %, &, etc.) can also be used to further increase the complexity of your password.

We have shared your comments with our management team for further review and consideration. Please continue to share your feedback with us as it is one of the best ways for us to improve.

Thank you for participating in the PC Plus program.

Bastaro I PC Plus Member Services

Troy: Hang on - are they saying they had customer feedback saying they should limit passwords to 8 characters?!

Banking (PC Financial) institutions require certain passwords because of old protocol. Same thing as getting upset over how insecure SMS is. We know. Who's going to pay to fix it?

I've worked on a banking system where as you say the backend system was an older

mainframe app with a limit of 8 characters for the password.

That however is no reason to have 8 character passwords on the web app that fronts the system. It's neither particularly difficult or expensive to have a modern password system over the front that doesn't have those restrictions.

Especially for a bank or financial institution it would be irresponsible not to.

My online bank enforces a password length of exactly 5 characters. That's right, five (as in "easily enumerable") characters. An online bank. WTF?

Epilogue

This post has actually been super useful when running my workshops. Time and time again, I'd talk about some crazy security thing I'd seen only to be met with scepticism from the audience; could anyone really do something as stupid as Troy is describing?! Here, hold my beer...

An opportunity I missed in this post was to refer to Poe's law:

"Poe's law is an adage of Internet culture stating that, without a clear indicator of the author's intent, it is impossible to create a parody of extreme views so obviously exaggerated that it cannot be mistaken by some readers for a sincere expression of the views being parodied."

Lars' tweet at the beginning of the post demonstrates this perfectly and Poe's law as it relates to infosec is truer today than ever. And that's one of the things I absolutely love about this industry; that it's constantly crazy enough for me to do a double-take and question whether the asserted security thing is true or not. I'll leave you with a perfect example of this from only just before the time of writing: a mate of mine called Ken Munro runs a pentest company in the UK (you'll read more about Ken when you get to the TicTocTrack blog post), and they recently discovered a vulnerability in a male chastity lock that could

enable an attacker to lock a guy's dick in place. How could this possibly be real?! Is it real? Yes. Shortly thereafter, Lorenzo from Vice writes a piece titled "Your Cock Is Mine Now:' Hacker Locks Internet-Connected Chastity Cage, Demands Ransom" which, of course, leads to much shock and awe amongst the community. Was it real? No, he was pranked by an Aussie comedian who went to some lengths to ensure he could fool the seasoned infosec journalist into not only believing the guy's wedding tackle was locked and ransomed, but that he also suffered an injury whilst trying to remove it with, uh, bolt cutters.

Like I said, I love this industry 😎

THE ETHICS OF RUNNING A DATA BREACH SERVICE

It was approaching 4 years of running HIBP and I was increasingly conscious of the profile it was gaining. This thing was everywhere; online articles, prime time news, me standing up in front of hundreds of people at conferences, it was just all over the place. The spotlight was well and truly on my little project and with that, came the criticisms.

One recurring theme was around whether it was irresponsible to allow anyone to search HIBP and find anyone else's exposure in data breaches. HIBP started out this way when it was a tiny project (by comparison) and that was just the default position I took. Over time though, it became much clearer that making data publicly searchable wasn't just a default position, it was a feature. There are many good reasons for this and since I was being increasingly interrogated about them, I needed a clearly articulated position on it. This is that blog post.

25 SEPTEMBER 2017

o matter how much anyone tries to sugar coat it, a service like <u>Have I</u> been pwned (HIBP) which deals with billions of records hacked out of other peoples' systems is always going to sit in a grey area. There are degrees, of course; at one end of the spectrum you have the likes of <u>Microsoft and Amazon using data breaches to better protect their customers' accounts</u>. At the other end, there's services like <u>the now defunct LeakedSource</u> who happily sold our personal data (including mine) to anyone willing to pay a few bucks for

it.

As far as intent goes, HIBP sits at the white end of the scale, as far to that extreme as I can possibly position it. It's one of many things I do in the security space alongside online training, conference talks, in-person workshops and of course writing this blog. All of these activities focus on how we can make security on the web better; increasing awareness, reducing the likelihood of a successful attack and minimising the damage if it happens. HIBP was always intended to amplify those efforts and indeed it has helped me do that enormously.

What I want to talk about here today is why I've made many of the decisions I have regarding the implementation of HIBP. This post hasn't been prompted by any single event, rather it seeks to address questions I regularly see coming up. I want to explain my thinking, explore why I've made many of the decisions I have and invite people to contribute comments with a hope of making it a more useful system for everyone.

The Accessibility of a Publicly Searchable System

The foremost question that comes up as it relates to privacy is "why make the system publicly searchable?" There are both human and technical reasons for this and I want to start with the former.

Returning an immediate answer to someone who literally asks the question "have I been pwned?" is enormously powerful. The immediacy of the response addresses a question that's clearly important to them at that very moment and from a user experience perspective, you simply cannot beat it.

The value in the UX of this model has significantly contributed to the growth of the service and as such, the awareness its raised. A great example is when you see someone take another person through the experience: "here, you just enter your email address and... whoa!" The penny suddenly drops that data breaches are a serious thing and thus begins the discussion about password strength and reuse, multi-step verification and other fundamental account management hygiene issues.

The fact that someone can search for someone who is not them is a double-edged sword; the privacy risk is obvious in that you may discover someone was in a particular data breach and then use that information to somehow disadvantage them. However, people also extensively use the service to help protect other people, for example by identifying exposed spouses, friends, relatives or even customers and advising them accordingly.

I heard a perfect example of this just the other day when speaking to a security bod in a bank. He explained how HIBP was used when communicating with customers who'd suffered an account takeover. By highlighting that they'd appeared in a breach such as LinkedIn, they are able to help the customer understand the potential source of the compromise. Without being able to publicly locate that customer in HIBP, it would be a much less feasible proposition for the bank.

I mitigate the risk of public discoverability adversely impacting someone by flagging certain breaches as "sensitive" and excluding them from publicly visible results. This concept came in when the Ashley Madison data hit and the only way to see if you're in that data breach (or any other that poses a higher risk of disadvantaging someone) is to receive an email to the searched address and click on a unique link (I'll come back to why I don't do that for all searches in a moment).

I've actually had many people suggest that it's ok to show the sensitive results I'm presently returning privately because the privacy of these individuals has already been compromised due to the original breach. I don't like this argument and the main reason is because I don't believe the act of someone else having illegally broken into another system means the victims of that breach should be

exposed further in ways that would likely disadvantage them. It's not the only time I've heard this, for example after launching <u>Pwned Passwords</u> last month a number of people said "you should just return email address and password pairs because their data is out there anyway". Shortly after that, I was told I'm "holding people hostage" by not providing the passwords for compromised email addresses. In fact, I had someone get quite irate about that after loading <u>the Onliner Spambot data</u> with the bloke in question then proceeding to claim that not disclosing it wasn't protecting anyone. Someone else suggested I was "too old fashioned and diplomatic". No! These all present a significantly greater risk for those individuals and for someone who himself is in HIBP a dozen times now, I'd be pretty upset if I saw any of this happening.

Searching by Email Verification is Fraught with Problems

This is the alternative I most frequently hear – "just email the results". There are many reasons why this is problematic and I've already touched on the first above: the UX is terrible. There's no immediate response and instead you're stuck waiting for an email to arrive. Now you may argue that a short wait is worth the trade-off, but there's much more to it than that.

HIBP gets shared and used constructively in all sorts of environments that depend on an immediate response. For example, it gets a huge amount of press and a search is regularly shown in news pieces. Many people (particularly in the infosec community) use it at conference talks and they're not about to go opening up their personal email to show a result.

But those are arguments in favour of accessibility and I appreciate not everyone will agree with them so let's move onto hard technical challenges and the first is delivering email. It's very hard. In all honesty, the single most difficult (and sometimes the single most expensive) part of running this service is delivering

mail and doing it reliably. Let's start there actually - here's <u>the cost of sending</u> 700k emails via SendGrid:

	SHARED IP	BOOST DELIVERABILITY WITH A DEDICATED IP	
Feature Highlight	Essentials	Pro	Premier
Monthly plan price	\$19.95	\$399.95	Custom
Emails/month included in plan	100,000	700,000	2.5+ million
Price per extra email	\$0.00075	\$0.00045	Custom
APIs and webhooks	~	~	~
DKIM oustomization	~	~	~
SendGrid Marketing Campaigns	~	~	~
Dedicated IP		~	~
Subuser management		~	~
Customer Success Manager			~
Prioritized support			~

Now fortunately, SendGrid helps support the project so I don't end up wearing that cost but you can see the problem. Let's just put the challenge of sending an email on every search in context for a moment: a few weeks ago, I had 2.8 million unique visitors in just one day after making the aforementioned 711 million record Onliner Spambot dump searchable. Each one of those people did at least one search and if I was to pay for that volume, here's what I'd be looking at:

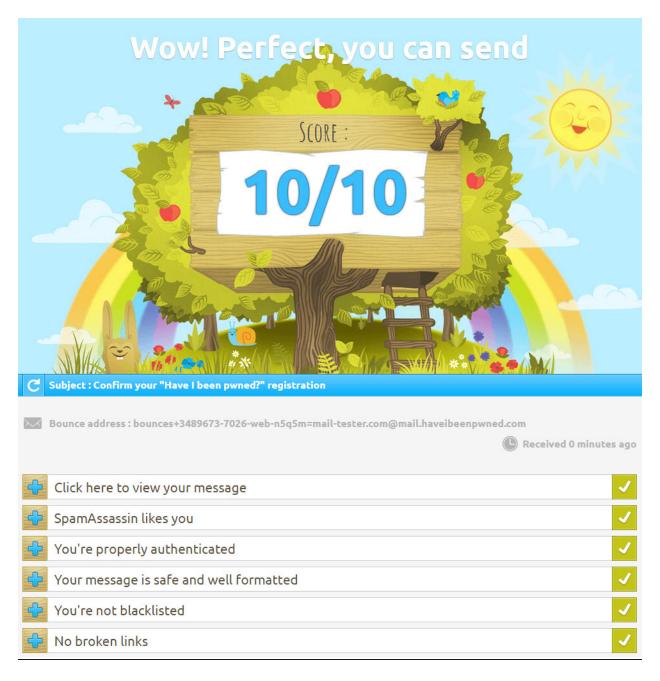
1. Estimate your needs

Send around 3,000,000 emails per month

Try for Free* or buy from \$1,174.95/mo

That's one day of traffic. I can't run a free service that way and I hate to think of the discussion I'd be having with my wife if I did! Now that was one *exceptional* day but even in low periods I'm still talking about many millions of visitors a month. As it is, I'm coming very close to maxing out my email allocation each month just from sending verification emails and notifications when I load breaches. And no, a cheaper service like <u>Amazon SES</u> is not a viable alternative, I've been down that path before and it was a debacle for many reasons plus would still get very pricey. (Incidentally, large volumes of emails in a spike often causes delivery to be throttled which would further compound the UX problem of people waiting for a search result to land.)

And then there's the deliverability problems. One of the single hardest challenges I have is reliably getting mail through to people's inbox. <u>Here's what my mail setup looks like in terms of spam friendliness</u>:



DKIM is good. SPF is good. I have a dedicated IP. I'm not on any black lists. Everything checks out fine yet consistently, I hear people say "your notification went to junk". I suspect it's due to the abnormal sending patterns of HIBP, namely that when I load a breach there's a sudden massive spike of emails sent but even then, it's only ever to HIBP subscribers who've successfully double-opted-in. So, think of what that would mean in terms of using email as the sole channel for sharing breach exposure: a heap of people are simply going to miss

out. They won't know they were exposed in a breach, they won't adapt their behaviour and for them, HIBP becomes useless.

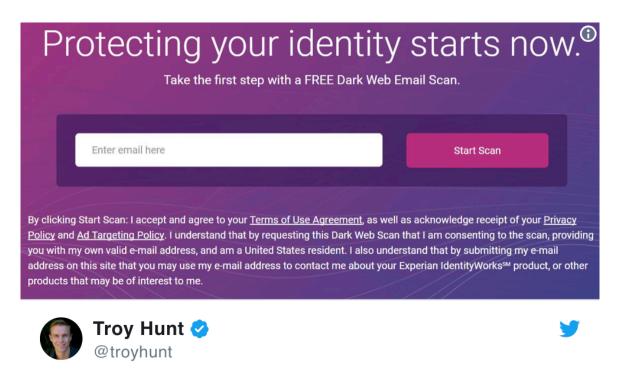
I've seen criticism from other services attempting to do similar things to HIBP based on the fact I'm not just sending emails to answer that "have I been pwned?" question. But they're at a very different stage of maturity and popularity and simply don't have these challenges – it'd be a lot easier if I was only sending hundreds of emails a day and not tens or sometimes even hundreds of thousands. They're also often well-funded and commercialise their visitors so you can see why they may not understand the unique challenges I face with HIBP.

In short, this is the best possible middle ground I can find. Not everyone agrees with it, but I hope that even the folks who don't can see it's a reasoned, well thought out conclusion.

Because I Don't Want Your Email Address

There are a number of different services out there which offer the ability to identify various places your data has been spread across the web. It's a similar deal to HIBP insofar as you enter an email address to begin the search, but many then promise to "get back to you" with results. Of course, during this time, they retain your address. How long do they retain it for? Well...

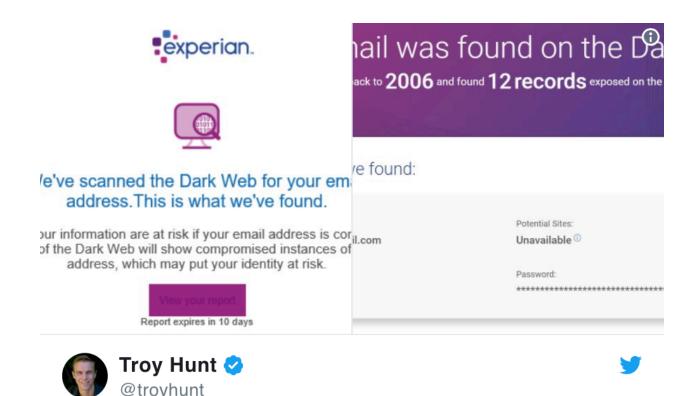
Someone directed me to <u>Experian's "Dark Web Email Scan" service</u> just recently. I had the feeling just from reading the front page that there was more going on than meets the eye so I took a look at the policies they link to and that (in theory) you must read and agree to before proceeding:



It's stunts like @Experian is pulling that erode trust in these companies: that's 21,494 words you need to agree to: experian.com/consumer-produ...

♡ 83 7:21 PM - Sep 8, 2017 · Melbourne, Victoria

Folks who've *actually read all this* have subsequently pointed out that as expected, providing your address in this way now opts you into all sorts of things you really don't want. In fact, I saw it myself first hand:



Geez the Experian "dark web search" is terrible: several days to get a result, useless info in the report and 2 subsequent spam mails since

○ 37 10:49 AM - Sep 13, 2017 · Gold Coast, Queensland

In other words, the service is a marketing funnel. The premise of "just leave us with your address and we'll get back to you" is often a thin disguise to build up a list of potential customers. Part of the beauty of HIBP returning results immediately is that the searched address never goes into a database. The only time this happens is when the user explicitly opts in to the notification service in which case I obviously need the address in order to contact them later should they appear in a new data breach. It's data minimisation to the fullest extent I can; I don't want anything I don't absolutely, positively need.

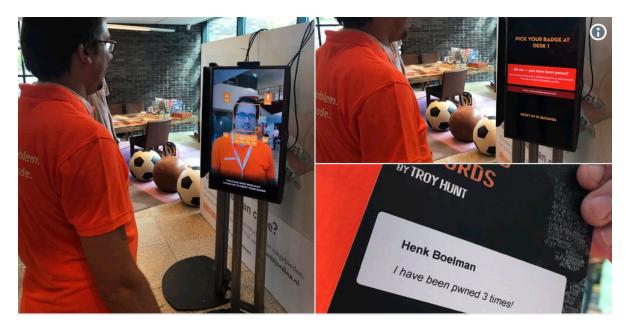
Incidentally, by Experian not explicitly identifying the site the breach occurred on it makes it *extremely* difficult for people to actually action the report. They're not the only ones - I've seen other services do this too - and it leaves the user

thinking "what the hell do I do now?!" I know this because it's precisely the feedback I had after loading the Onliner Spambot data I mentioned earlier, the difference being that I simply didn't know with any degree of confidence where that data originated from. But when I do, I tell people - it's just the right thing to do.

The API is an important Part of the Ecosystem

One of the best things I did very early on in terms of making the service accessible to a broad range of people was to <u>publish an API</u>. In additional to that, <u>I list a number of the consumers of the service</u> and they've done some great things with it. There are many other very good use cases you won't see publicly listed and that I can't talk about here, but you can imagine the types of positive implementations ingenious people have come up with.

In many cases, the API has enabled people to do great things for awareness. For example, this implementation at a user group I spoke at in the Netherlands recently (and yes, opt-in was optional):







So @ordina set up facial recognition via photos uploaded at registration & checked people against @haveibeenpwned on arrival #TroyAtOrdina

♡ 159 6:15 AM - Jul 5, 2017 · Nieuwegein, Nederland

The very nature of having an API that can search breaches in this fashion means the data has to be publicly searchable. Even if I put API keys on the thing, I'd then have the challenge of working out who I can issue them to then policing their use of the service. For all the reasons APIs make sense for other software projects, they make sense for HIBP.

Now, having said all that, the API has had to evolve over time. Last year I introduced a rate limit after seeing usage patterns that were not in keeping with ethical use of the service. As a result, one IP can now only make a request every 1.5 seconds and anything over that is blocked. Keep it up and the IP is presented with a JavaScript challenge at Cloudflare for 24 hours. Yes, you can still run a lot of searches but instead of 40k a minute as I was often seeing from a single IP, we're down to 40. In other words, the worst-case scenario is only one one-

thousandth of what it previously was. What that's done is forced those seeking to abuse the system to seek the data out from other places as the effectiveness of using HIBP has plummeted.

Like many of the decisions I've made to protect individuals who end up in HIBP, this one has also garnered me criticism. Very often I feel like I'm damned if I do and I'm damned if I don't; some people were unhappy in this case because it made some of the things they used to do suddenly infeasible. Yes, it slashed malicious use but you can also see how it could impact legitimate use of the API too. I'm never going to be able to make everyone happy with these decisions, I just have to do my best and continue trying to strike the right balance.

I'm Still Adamant About Not Sharing Passwords Attached to Email Addresses

A perfect example of where I simply don't see eye to eye with some folks is sharing passwords attached to email addresses. I've maintained since day 1 that this poses many risks and indeed there are many logistical problems with actually doing this, not least of which is the increasing use of stronger hashing algorithms in the source data breaches.

Not everyone has the same tolerance to risk in this regard. I mentioned earlier how some especially shady services will provide your personal data to anyone else willing to pay; passwords, birth dates, sexualities - it's all up for grabs. Others will email either the full password or a masked portion of it, both of which significantly increase the risk to the owner of that password should that email be obtained by a nefarious party.

I've tried to tackle the gap between providing a full set of credentials and only the email address by <u>launching the Pwned Passwords service last month</u>. Whilst the primary motivation here was to provide organisations with a means of

identifying at-risk passwords during signup, it also helps individuals directly impacted by data breaches; find both your email address and a password you've used before on HIBP and that's a pretty solid sign you want to revisit your account management hygiene.

At the end of the day, no matter how well I was to implement a solution that attached email addresses to other classes of personal data, there's simply no arguing with the basic premise of I cannot lose what I do not have. I have to feel comfortable with the balance I strike in terms of how I handle this data and at present, that means not putting it online.

There Are Still a Lot of Personal Judgement Calls

I've been asked a few times now what the process for flagging a breach as sensitive is and the answer is simply this: I make a personal judgement call. I have to look at the nature of the service and question what the impact would be if HIBP was used as a vector to discover if someone has an account on that site. I don't always get this right; I didn't originally flag the Fur Affinity breach as sensitive because I didn't understand how furries can be perceive until someone explained it to me. (For the curious, it's the sexual aspects of furries that came as news to me.)

HIBP is a constant series of judgement calls when it comes to the ethics of running the service. The data I should and should not load is another example. I didn't load the Australian Red Cross Blood Service breach because we managed to clean up all known copies of it (there are multiple reasons why I'm confident in that statement) and they committed to promptly notifying all impacted parties which they summarily did. I removed the VTech data breach because it gave parents peace of mind that data relating to kids was removed from all known locations. In both those cases, it was a judgement call made entirely of

my own free volition; there were no threats of any kind, it was just the right thing to do.

HIBP is not about trying to maximise the data in the system, it's about helping people and organisations deal with serious criminal acts. Frankly, the best possible outcome would be for there to be no more breaches to load. This is what all my courses, workshops, conference talks and indeed hundreds of blog posts are trying to drive us towards – fixing the problems that have led to data breach search services being a thing in the first place. Not everyone has those same motives though, and that's leading us to some pretty shady practices.

The "No Shady Practices" Rule

As I said in the intro, there's no sugar-coating the fact that handling data breaches is always going to sit in a grey area. This makes it enormously important that every possible measure is taken to avoid any behaviour whatsoever that could be construed as shady. It probably shouldn't surprise anyone, but this is not a broadly held belief amongst those dealing with this class of data.

I mentioned LeakedSource earlier on; there are still multiple sites following the same business model of "give us a few bucks and we'll give you other people's data". There's a total disregard not just for the privacy of people like you and I, but for the impact it can then have on our lives. People bought access to my own data – I know this because someone once sent it to me! Many of these services operate with impunity under the assumption that they're anonymous; great lengths are gone to in order to obfuscate and shield the identity of the operators although as we saw with Leaked Source, anonymity can be fleeting.

There are also multiple organisations paying for data breaches. What this leads to is criminal incentivisation; rewarding someone for breaking into a system and pilfering the data in no way improves the very problem these services set out to address. Mind you, the argument could be made that the purpose these services primarily serve is to be profitable and viewed in that light, paying for data and then charging for access to it probably makes sense from an ROI perspective. I've never paid for data and I never intend to and yes, that means that it sometimes takes longer for it to appear on HIBP, but it's the right thing to do.

Ambulance chasing is another behaviour that's well and truly into the dark end of shady. I recently had a bunch of people contact me after an organisation emailed them to advise that addresses from their company were found in a breach. Then I watched just last month as someone representing another org hijacked Twitter threads mentioning HIBP in order to promote their own service (I then had to explain what was wrong with this practice, something I later highlighted in another thread). In all these cases, financial incentive either from directly monetising the service itself or indirectly promoting other services associated to the organisation appear to be the driver for shady practices.

We should all be beyond reproach when handling this data.

Summary

Being completely honest, it would have to be less than one in one thousand pieces of feedback I get that are critical or even the least bit concerned about the HIBP model as it stands today. It's a very rare thing and that may make you wonder why I even bothered writing this in the first place, but the truth is that it helped me get a few things straight in my own head whilst also providing a reference point for those who *do* express genuine concern.

HIBP remains a service that first and foremost serves to further ethical objectives. This primarily means raising awareness of the impact data breaches are having and helping those of us that have been stung by them to recover from the event. Even as I've built out commercial services for organisations that have requested them, you won't find a single reference to this on this site; there's no

"products" or "pricing", no up-sell, no financial model for consumers, no withholding of information in an attempt to commercialise it, no shitty terms and conditions that you have to read before searching and not even any advertising or sponsorship. All of this is simply because I don't want *anything* detracting from that original objective I set forth.

I'll close this post out by saying that there will almost certainly be changes to this in the future. Indeed, it's constantly changed already; sensitive breaches, rate limits and the removal of the pastes listing are all examples of where I've stepped back, looked at the system and thought "this needs to be done better". Very often, that decision has come from community feedback and I'd like to welcome more of that in the comments below. Thank you for reading.

Comments

A good read Troy; I have to say that I am often amazed at how calmly you respond to the "one in one thousand pieces of feedback" considering how idiotic some of the comments/ tweets I have seen are. I think I would be shouting a bit more if it was me being hit for providing a free service like this! Beer is on me if you are ever in the West of England (or I ever manage to make it to a talk).

Troy: Admittedly, some of the comments I've received have really tested my patience. The thing is though, you never know whether the person is just having a bad day, is just messing with you or is genuinely upset. One cranky enough person can make life really difficult and I'd far prefer to just diffuse the situation rather than let it escalate into something more than it needs to be.

Some of the worst of it is just funny though, you must get a mental picture of someone frothing at the mouth, hammering away at a keyboard, without the faintest clue what they

are really talking about?

[... that's not a confession, by the way!]

Troy: Yep, in some cases I'm pretty sure that's precisely what's going on...

Some of the comments I have read seem to fall into the "I have really strong opinions on [insert matter here] and must voice them; but I will not listen to reason because I am correct... always".

In the world of internet comments defusing the situation is often the best; if not always the least painful!

Simply because you mentioned it, I'm curious what issues do you feel SES face compared to Sengrid? We've been using Sengrid since Mandrill ditched SaaS providers and we're looking at SES for a few reasons.

Troy: I should have seen this coming:)

The biggy was that you need a very high reputation to send with them. As soon as it took a little hit (i.e. I loaded a new breach which resulted in a bunch of emails sending at once), I'd get a warning and a demand to improve the rating. This was exceptionally difficult because I was only emailing people who'd double-opted-in anyway.

From memory, there also wasn't the ability to review outbound logs. One of the things I find myself doing time and time again is jumping into the SendGrid UI and checking if an address is bouncing after they advise email isn't landing properly.

Also, there was no dedicated IP when I was originally using the service (about 18 months ago), so you have problems with reputation due to other senders.

They offer it now, but suggest "a sustained and consistent sending pattern and a minimum daily sending volume of 175,000 emails per day (on average)".

Even with SendGrid and everything configured perfectly, I'm suffering. I recently needed help from friends at Microsoft to get off an Office 365 blacklist due to HIBP being flagged by a third party reputation service. They couldn't tell me which one or why I was on there.

Successfully delivering email for a service with usage patterns like this is a constant challenge and even folks who are very good at managing email reputation have been left scratching their heads.

Would it be possible for you split up the send and push it out from multiple servers? Typically when we were sending huge email sends we would use 10-15 servers (a provisioned VM) and split it across the lot. Your setup is quite different but maybe the same principle could be applied

Troy: Then I'd need additional IP addresses from SendGrid which comes at additional cost. Besides, 99.x% of the time there's no problem with the volume of content that's being sent, it's only when loading a very large breach that problems emerge.

Epilogue

What's really bugged me about this topic in the years that have passed since I wrote about it is someone (let's just call him "Barry", for the sake of example) popping up with very strongly held opinions about the public searchability of data but with little basis behind them. Barry sees this thin veneer of an online system, observes the behaviour then sends me a ranty email about how unhappy he is. I'm sitting there thinking "have you got any idea of how much

time I spend thinking about this shit?!" I can understand that we all form our own opinions based on the knowledge and experiences we each have but without sounding too obnoxious about it, mine is rather more, uh, "extensive" than Barry's.

Let me share an experience that came 18 months after writing this blog post. I did a talk at the NYPD Cyber Intelligence and Counterterrorism Conference in New York and after the keynote ended, a bunch of law enforcement officers came by for a chat. (The good kind of chat!) They thanked me for making HIBP available as they found it a really useful tool for both their internal training and community outreach programs. Being able to sit there with someone, plug their email address in and get an immediate result was enormously valuable to them. Barry didn't know that, but then how would he?

Years on from writing this post, I'm stauncher than ever that the balance is right. That doesn't mean the feature is never abused, I accept that's a possibility, but what it means is that all the evidence I have heavily weights the decision towards ease of searching above and beyond verification of address ownership. It all comes back to a key principle I've stuck by since day one: data breaches happen, now what's the best thing we can do in the wake of them? HIBP is my answer.

HERE'S HOW I DECIDE WHAT I ENDORSE AND HOW I ENSURE TRANSPARENCY

One of the unexpected upsides of building out a public profile has been the opportunity to take something that was a passionate hobby and turn it into an actual career. The only way that can happen is if you can monetise it because at the very least, you've gotta pay the bills. But with a growing public profile and growing influence came an all-new challenge – how do I deal with people who want to buy that influence? Don't get me wrong – this is a good problem to have – but it is a problem insofar as whilst endorsements can be attractive (and we're talking everything from free tech to actual cash payments), they also run the risk of damaging your credibility. I'd built out a brand as a transparent and trustworthy voice in the industry and that's something that took years to achieve, I could blow all that very quickly by attaching myself to the wrong organisation for the wrong reasons.

As with many of my blog posts, this one was as much about getting things straight in my own head as it was about communicating publicly. Where did my own moral compass sit on this topic? Was I clear about this or just winging it as I went along? In truth I think it was probably a bit of both until I sat down and thought about it carefully. Today, this post still holds very true, and I take the same view with all my professional relationships.

03 OCTOBER 2017

ne of the by-products of an increasingly public profile is that companies want you to promote their things. You see this all the time in all walks of life whether it be product placement in movies, celebs sponsored by car companies or indeed the sponsor banner you see at the top of this blog. These companies benefit from the exposure granted to them by individuals with influence.

The flip side is that the allure or money or free goods can taint the impartiality of said individual. For example, in the wake of the Sony Pictures hack we learned that Kevin Hart was paid a couple of million bucks to tweet Sony's messages. More recently, there was news that the Kardashian family wasn't properly disclosing paid endorsements on Instagram. Now I don't exactly have those sorts of levels of celebrity status but it did get me thinking - how do I decide what stuff I attach my name to?

This post makes that position clear. It's necessary because I'm increasingly asked about it and indeed, often challenged by people who believe I may lack impartiality. Here's how I handle it.

I'm Not Endorsing Anything I Wouldn't Use Myself

The very best relationships I have with companies are the ones where they've approached me after my own independent endorsement of what they do. For example, I have the Microsoft MVP and Regional Director roles (no, for the millionth time that doesn't mean I work for Microsoft!) and they came after extensively using their products in a professional capacity for a decade or more. Not just using them either, but writing about them at length and being a community influencer. My series on The OWASP Top 10 for .NET Developers was the catalyst for that relationship and the MVP award came another year after I began that journey.

Lenovo is another great example. I've used ThinkPads since the 90's, originally when they were IBM then into the Lenovo era that began in '05. That was all Pfizer bought us when I worked there - 14 year's worth of ThinkPads! They were great machines so when I needed a new one 4 years ago now, I went out and spent my own hard-earned cash on one:

Congratulations!

You have committed to buy:

LENOVO THINKPAD W540 i7-4800QM, 3K, IPS, 2880 x 1620 DISPLAY, 32GB, 512GB SSD $\underline{151436818344}$

Be sure to check out my complete list of items for sale in my store.

VIEW OUR STORE SEARCH EBAY

These days, I'm on their Insiders program and they send me machines, but they're precisely the sorts of machines I'd buy myself anyway. The P50 I received last year is without doubt the best machine I've ever had, regardless of how I acquired it.

Then there's Ubiquiti and if you've been watching, I've said quite a lot about them recently. That all started because I went and spent a couple of grand of my own to finally fix my dodgy wifi. And it's awesome! That blog post was written before ever speaking to anyone there or receiving so much as a free sticker from them.

Now I never expected this to happen to the degree it has, but after writing that blog it turns out that a lot of people went and bought Ubiquiti bits. What makes me especially pleased about the results in that tweet search is how happy everyone is - people *love* the gear and that's really important to me

independently of any commercial interests I have. Which brings me to the next point:

I Always Make It Crystal Clear if I'm Financially Incentivised

Continuing the Ubiquiti topic, regardless of how independently endorsed I am in a product, it's enormously important that I disclose when I've been financially incentivised. For example, when I recently wrote the course Ubiquiti has now put out for free I said this:

I want to be clear that this is a commercial course (they've paid me for my time)

Or when I managed to get 7 aesthetically faulty "factory second" in-wall units that I put into my brother's house in a ground up build:

Functionally they were perfect, but they weren't yet 100% happy with the fitment of the covers. But if I wasn't the fussy type, how many did I need and would I like them to send me over a box of near perfect ones for free? 7, and yes please:)

This is really important context and whilst I can't always fit disclosures into say, a single tweet, I make sure that at every opportunity I'm clear about the relationship. Frankly, I don't think this takes anything away from the value the company in question gets out of the relationship (they still get the same exposure) and if I *didn't* properly disclose, there's a very real chance it would take something away from me, namely my independence and authenticity.

All of that said, it's not like I say "yes" to any company that pops up and wants to throw money or product my way either. For example:

I Say "No"

Just recently, I had a company you know approach me to write for them. I won't name them here, but let's just say they rhyme with an Irishman who doesn't move too fast (some of you will get it). They wanted me to write some content on a commercial basis but the historical reputation of the company just didn't sit well with me. They've done some shitty things in the past and whilst in more recent years they've clearly tried to turn that around, the bad memories are still just too fresh.

When any of us attach our names to another company whether that be by writing about them, writing for them or even just publicly using their products, a bit of us rubs off on them and a bit of them rubs off on us. It can be a delicate balance and in a case like this particular one, an argument could be made that a positive security influence from someone such as myself is in the industry's best interests. But I didn't feel I could really achieve that through writing alone and that the net result of that relationship would be negative for me.

The other day I had a company contact me asking "if we can cooperate with you" which as it turned out, was code for "can you please talk about our things if we give them to you for free". I'm pretty uncomfortable with this premise - it just doesn't sit well with me. The exception of course is if it's something I'm already endorsed in which is as I explained earlier on.

Same again for blog sponsorship. I rolled that model out in September last year and it's been *fantastic!* But I've said "no" on multiple occasions because I didn't agree with the philosophy of the company wanting space on my site. Now of course, sponsor messages are a different proposition to me directly endorsing someone's laptops or wifi gear; I haven't personally used many of the products my blog sponsors are selling (with a few notable exceptions), but they're brands and names I'm happy to have occupying that bar for one-week slots at a time.

The point is that saying "no" is ok. But when I do say "yes", there's another really important aspect of every endorsement:

I Maintain Full Independence

I didn't like Lenovo's Superfish, but their laptops are the best I've ever used. I've subsequently leveraged my relationship with them on multiple occasions to talk specifically about Superfish and what I believe they need to do better.

I don't like Microsoft's lack of support for browser security standards such as SRI and HPKP. So, I use the influence I have with them to express why that's important and why I believe it needs to change. (I'm still pushing them for first class Let's Encrypt in the Azure App Service too, by the way.)

The point is that every organisation has strengths and weaknesses and having a commercial relationship shouldn't mean only talking about the former and neglecting the latter. That adversely impacts credibility and quite rightly, makes people question your independence.

I can quite honestly say that none of the organisations I've worked with have ever had anything to say when I've publicly talked about what I don't like. I've also never had any push-back when I've explained why I don't want to do things like hashtag tweets or includes logos in email footers. I've always explained to them that the value I represent is that people trust my transparency and candour; anything that jeopardises that would be negative for all involved.

Ultimately, I approach every relationship with one simple objective, and it's this:

You Shouldn't Be Able to Tell the Difference

I mean you shouldn't be able to tell the difference between a product I endorse because I'm incentivised to do so versus one I just like, short of the disclosure I mentioned earlier, of course. My behaviour shouldn't change just because there's money or product changing hands.

For example, I'm frequently *very* vocally supportive of both the <u>1Password</u> <u>password manager</u> and <u>Freedome VPN</u>. I use them both daily, I've written about them both and I constantly recommend them to anyone who asks. I've never received either product for free (I've paid retail prices for both for years), and I've never been paid to endorse either of them. I have contacts at both companies I've spoken to on various issues multiple times in the past, but any sort of advocacy position is simply not a topic that's ever come up.

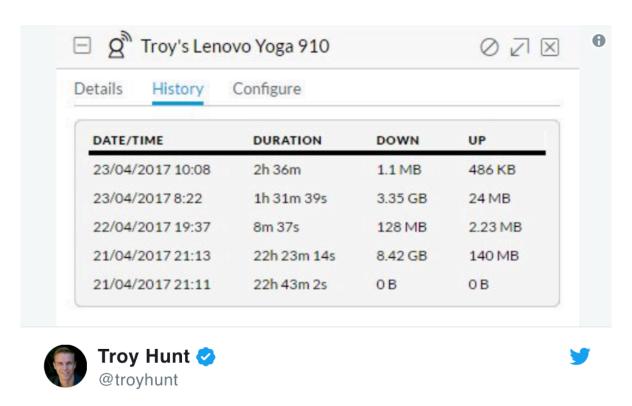
So, to the point of the title, the way I talk about 1Password and Freedome should be indistinguishable to the way I talk about Lenovo and Ubiquiti in terms of how I endorse them. It's just the right thing to do.

Now, having said all that, no matter how hard I try to get all this right, I'm never going to keep everyone happy all of the time:

People Will Still Complain Anyway

Let me give you an example and it's not to throw this guy under the proverbial bus (he was actually very cool in his later responses), but rather it's to highlight the ongoing challenge of finding the right balance.

Recently, Lenovo sent me a new machine, a Yoga 910 in this case. This was running Win 10 out of the box but I was curious - how much data does a *brandnew* machine still need to pull down from the web just to get up to current patch levels? Nothing to do with the fact I had a free machine from Lenovo, just genuine geeky curiosity. So, I logged it and tweeted this:



Brand new machine out of the box only used to take Windows updates and new drivers: 12GB of downloads thank you very much! (stats via @ubnt)

○ 34 5:01 PM - Apr 22, 2017

Seems fair, right? Interesting even? Not everyone thought so:

Sorry Troy, I've had to Unfollow you. Can't handle anymore of your Ubnt promos being thrown down my Twitter feed.

If anything, I thought I'd get a harsh comment or two for having received a freebie from Lenovo! The only reason I added the @ubnt reference was because if I *didn't*, I'd get a barrage of responses along the lines of "where did you get those stats from".

Admittedly, when I saw this I was pissed. I had a really good reason for mentioning Ubiquiti and it had nothing to do with promoting them. But instead of responding as I felt inclined to, I took a different approach:



Because that's ultimately what it boils down to, right? If you're following someone and their signal to noise ratio moves outside your comfort zone then you simply don't follow them anymore. Ok, most people probably wouldn't explicitly broadcast their intent to unfollow, but obviously the guy was frustrated and he wanted to vent that. I added a smiley face in my reply because I didn't want it to appear condescending, and clearly it didn't, because I got this back a few minutes later:

> Gosh, I almost want to (re) follow you, simply because this is the most sensible response I could have got.

At the time of writing, he still follows me $\stackrel{\textstyle \longleftarrow}{}$.

The point is that there's always going to be people that fundamentally disagree with your point of view and once you add a variable that is perceived to impact your impartiality, that disagreement is amplified. That's why I'm so cautious with everything I attach my name to and indeed, that's why I've written this post - to explain my thought process. I hope this helps explain things to all the people I'm sure I'll direct here in the future.

Epilogue

With the passage of time, this blog post has only served to reinforce itself and I'll give 3 good examples of why: a new relationship I've established, an old relationship I've terminated and a recent relationship I threatened to pull the

0

pin on.

The new one first and in late 2020, I became a strategic advisor for NordVPN. They made a good product, but their marketing people were sometimes prone to do what marketing people tend to do and embellish beyond what you could reasonably expect the product to do. For a company building a brand on trust and transparency, messaging that adversely impacted trust and transparency wasn't ideal. I had to give a lot of careful consideration to how the relationship would impact my own personal brand, ultimately deciding that by setting clear expectations of what I was there to do and leaving a clear exit if I wasn't able to achieve it, it would be a good relationship to have. And it has been, with the overwhelming majority of feedback being enthusiastic about the positive influence I could have.

The relationship I terminated was with a commercial certificate authority that had blog sponsorship. Now I'm a big proponent of free certs via the likes of Let's Encrypt so on the face of it, the very premise of paying for certs seems like it's not aligned with my own values. But there are cases where it makes sense, for example there are organisations with non-negotiable requirements to have support agreements with the CA, something you're not going to get for free. So, I carefully considered this sponsor and then brought them on board, being careful to ensure their sponsor message always aligned with my own personal views. What eventually brought it all undone was their views on both extended validation certificates (which were becoming increasingly useless), and their derogatory messaging around Let's Encrypt. I simply wasn't comfortable with their position, so I terminated the relationship and flagged them as "do not allow future sponsorships" (I literally have a column for this in my little CRM).

Most recently I had another sponsor who took umbrage with a tweet mentioning them during their sponsorship period. I'm not going to name them or provide enough info to figure out who it was because we ironed it all out, but it went like this: They reached out and asked me to delete the tweet because it wasn't aligned with their "brand and values". They also asked for any tweet mentioning them to have their input first so that they could contribute to "crafting it". I'll share snippets of my actual response because it gives a good insight into how much attention I apply to aligning my private messaging with my public messaging:

"This is not how sponsorship of my blog works; it's not an opportunity for the sponsor to define my tone or the way I communicate with my audience, nor is it an opportunity to request my messaging be censored."

"My brand and values are centred around transparency and authenticity and if your desire is to modify that in any way, [company] is not a good fit."

I was pissed. It was the first time ever I'd been asked to modify my messaging in this fashion, and it just wasn't going to happen. I offered to refund the sponsorship in full and immediately terminate the relationship, something I was frankly almost itching to do after their communication. They said they'd go away and discuss after which they came back and reconsidered their position.

What's really apparent to me today is both the power and responsibility that comes with having a marketable public profile. Now believe me - nobody finds it stranger than I do that I'm in this position - but the platform I've built has become enormously valuable. Increasingly, I need to be conscious of everything I say and do online as it comes under more scrutiny and has greater impact. But conversely, the influence is giving me a lot more power over situations like in that last example where the sponsor wanted to control my messaging; I imagined them sitting around the proverbial boardroom table discussing what the fallout would be if I was to pull their sponsorship and, god forbid, talk publicly about why I pulled it. That's kinda cool

HERE'S WHAT I'M TELLING CONGRESS ABOUT DATA BREACHES

"Hi, this is the US Congress, would you be cool to come and testify about data breaches?"

"Uh, you know I'm Australian, right?"

"Yeah, no problems, it'll be fine"

That's pretty much how it went. Following that was some to-ing and fro-ing about logistics, namely that I'd need to get over to literally the other side of the world (Washington DC), they couldn't pay me anything (not unexpected) and that I'd need to pay my own way there (unexpected). Actually, it was kind of expected insofar as expert witnesses need to be entirely independent and can't be incentivised which I can understand, but wow, that's a big trip. They suggested I reach out to my employer to pay for the trip because it'd be great exposure for them, but of course I was independent so that wasn't going to work. I decided that this would be an absolutely pivotal milestone in my career and that, if necessary, I'd just pay for the trip out of my own pocket. Fortunately, Pluralsight came to the rescue and picked up the tickets in exchange for me doing some events with them whilst I was there (it involved driving Porsches around a racetrack in Atlanta during a customer event, a sacrifice I was willing to make).

HIBP had always been a community-centric effort so I wanted to go to Congress with community-driven content, hence this blog post and the one I refer to at the beginning that preceded it. And I got some great stuff out of it too, things I still use today (the comment you'll read about data being a liability comes immediately to mind). I wrote this at a time that was equal parts exciting and daunting, having no idea what to expect once I turned up in DC. But hey,

data breaches were my thing, and I was confident talking about it, didn't matter if it was a little local user group or Congress. Breaches are breaches and just like I would at a conference talk, I carefully prepared, rehearsed then jetted off.

30 NOVEMBER 2017

ast week <u>I wrote about my upcoming congressional testimony</u> and wow-you guys are awesome! Seriously, the feedback there was absolutely sensational and it's helped shape what I'll be saying to the US Congress, including lifting specific wording and phrases provided by some of you. Thank you!

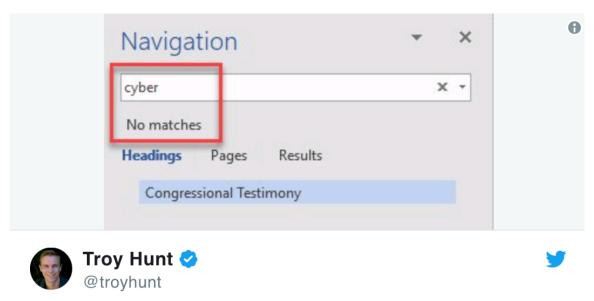
As I explained in that first blog post, I'm required to submit a written testimony 48 hours in advance of the event. That testimony is now <u>publicly accessible</u> and reproduced below.

Do keep in mind that the context here is the impact on identity verification in "a post-breach world".

My task is to ensure that the folks at the hearing understand how prevalent breaches are, how broadly they're distributed and the resultant impact on identity verification via knowledge-based authentication. I've had some great suggestions around tackling the root cause of data breaches and I'd love to have another opportunity in the future to talk about that, but it goes beyond the specific focus of this hearing. That said, who knows what I'll be asked by congressmen and congresswomen on the day and they may well question what can be done to combat the alarming rise in these incidents. I've now got a lot of great references on hand to go to should that happen so once again, thank you!

Below is the written testimony which has now been submitted and cannot be changed. (Incidentally, there were some formal requirements such as the 1-page summary in the opening of the document.) On Thursday morning at 10:15 DC time, I'll read my oral presentation which is a 5-minute distilled version of the

testimony below. I'll reproduce that in another blog post after the hearing as well as linking through to a recording of the event. In writing both of these, I've spent quite a bit of time watching previous hearings including Securing Consumers' Credit Data in the Age of Digital Commerce which featured Bruce Schneier (his written testimony is also worth a read). This has helped me pitch it at what I believe is the right level and I've reflected many of the terms and phrases I've heard from the folks on the committee. As such, you'll find this somewhat different to a lot of my usual writing as it's intended for a very different audience with a very different purpose. That said, I'm certainly not selling out on the things that are important to me!



Finalising my congressional testimony on data breaches. Big on facts, small on buzzwords

○ 620 2:46 PM - Nov 25, 2017

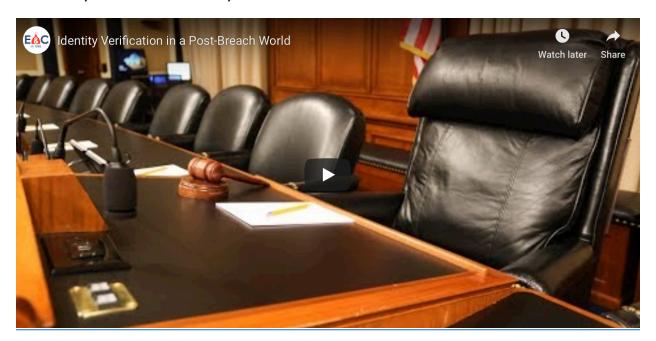
(Because someone will point it out if I don't mention it, yes, there is one "cyber" below which reflects Pluralsight's wording around their audience but *does not* appear in my oral presentation!)

Before you read this testimony, let me share one thing that's particularly noteworthy one week on: I wrote about the prevalence of old data breaches *before* the Uber news broke and before I was passed the imgur data. I wrote

about "unknown unknowns" before those breaches became "knowns" and whilst I'm not naming names in the testimony, I'm sure you can see significance of the timing, especially given the way the Uber situation was handled.

If you're around DC and want to come along, the notice of the hearing has all the info you should need (please come say hi if you do). If you can't be there but would like to tune in on the day, the hearing notice states that it will be available via webcast at energycommerce.house.gov. At the time of writing, there's a YouTube video sitting on the hearing page stating it will go live at the scheduled hearing start time. I've also embedded it below for convenience sake:

Here's my written testimony in full:



Statement of Troy Hunt

For the House Committee on Energy and Commerce

"Identity Verification in a Post-Breach World"

Summary

- Data breaches occur via a variety of different "vectors" including malicious activity by attackers exploiting vulnerabilities, misconfiguration on behalf of system owners and software products intentionally exposing data by design.
- There is frequently a long lead-time (sometimes many years) between a data breach and the service owner (and those in the breach) learning of the incident. We have no idea of how many incidents have already occurred but are yet to come to light.
- The industry has created a "perfect storm" for data exposure. The rapid emergence of cheap, easily accessible cloud services has accelerated the growth of other online services collecting data. Further to that, the rapidly emerging "Internet of Things" is enabling us to digitise all new classes of information thus exposing them to the risk of a data breach.
- An attitude of "data maximisation" is causing services to request extensive personal information well beyond the scope of what is needed to provide that service. That data is usually then retained for perpetuity thus adding to an individual's overall risk.
- Lack of accountability means that even in the wake of serious breaches, very little changes in the industry and we continually see other organisations repeat the same mistakes as their peers.
- Data breaches are redistributed extensively. There's an active trading scene exchanging data both for monetary gain and simply as a hobby; people

collect (and thus replicate) breaches.

- Many of the personal data attributes exposed in breaches cannot be changed once in the public domain, nor can these breaches be "scrubbed" from the internet once circulating.
- Even without data breaches, we're willingly exposing a huge amount of personal information publicly via platforms such as social media.
- The prevalence with which our personal data is exposed has a fundamental impact on the viability of knowledge based authentication. Knowledge which was once personal and could be relied upon to verify an individual's identity, is now frequently public knowledge.

Opening

Vice Chairman Griffith, Ranking Member DeGette, and distinguished Members of the House Energy and Commerce Committee, thank you for the opportunity to testify.

My name is Troy Hunt. I'm an independent Australian Information Security Author and Instructor for Pluralsight, an online learning platform for technology and cybersecurity professionals. I'm commissioned on a course-by-course basis to create training material that has been viewed by hundreds of thousands of students over the last 5 years. I'm also a Microsoft Regional Director (RD) and Most Valuable Professional (MVP), both titles of recognition rather than permanent roles. I've been building software for the web since 1995 and specialising in online security since 2010.

Of particular relevance to this testimony is my experience running the data breach notification service known as Have I Been Pwned (HIBP). As a security researcher, in my analysis of data breaches I found that few people were aware of their total exposure via these incidents. More specifically, I found that many people were unaware of their exposure across *multiple* incidents (one person appearing in more than 1 data breach) and indeed many people were unaware of *any* exposure whatsoever. In December of 2013, I launched HIBP as a freely accessible service to help people understand their exposure. Over the last 4 years, the volume of data in the service has grown to cover more than 250 separate incidents and over 4.8 billion records. What follows are insights drawn largely from running this service including the interactions I've had with companies that have been breached, those who have had their personal data exposed (myself included) and law enforcement in various jurisdictions around the world.

Data Breach Vectors

Data breaches have become a fact of modern digital life. Our desire to convert every aspect of our beings into electronic records has delivered both wonderful societal advances and unprecedented privacy risks. It's an unfortunate yet unavoidable reality that the two are inextricably linked and what follows describes the risks we are now facing as a result.

The term "data breach" is used broadly to refer to many different discrete vectors by which data is exposed to unauthorised parties. Some are as a result of malicious intent, some occur due to unintentional errors and yet others are inadvertent by-products of software design; they're "features", if you will.

Malicious incidents are the events we immediately associate with the term "data breach". In this case, a "threat actor" has deliberately set out to gain unauthorised access to a protected system, often with the intention of causing harm to the organisation and their subscribers. We frequently see successful attacks mounted through exploitation of very well-known vulnerabilities with

equally well-known defences. They exploit flaws in our software design, our security measures and indeed our human processes. They may be as sophisticated as leveraging previously unknown flaws or "zero days", yet they're frequently as simple as exploiting basic human shortcomings such as our propensity to choose poor passwords (and then to regularly reuse them across multiple services).

Especially in recent years with the growing ubiquity of easily accessible cloud services, data breaches often take the form of unintentionally exposed data. The ease today with which a publicly facing service can be provisioned and large volumes of data published to it is unprecedented – it can take mere minutes. Equally unprecedented is the simplicity with which an otherwise secure environment can be exposed to the masses; a single firewall setting or a simple access control change performed in mere seconds is all it takes.

The very design of some online services predisposes them to revealing large volumes of data about their subscriber base. Particularly in systems intended to make people discoverable such as social media or dating sites, we've seen many precedents of large volumes of publicly accessible information collated in an automated fashion in order to build a rich dataset. Some may be reluctant to even call this a "data breach", yet the end result is largely consistent with the previous two examples of malicious intent and unintentionally disclosed data.

We Often Don't Know Until Years Later

We simply have no idea of the scale of data that has been breached. We can measure what we know and conclude that there's an alarmingly large amount of personal information having been exposed, but it's the extent of the "unknown unknowns" that is particularly worrying.

Increasingly, we're realising the significance of the problem. During 2016 and 2017 in particular, we saw many incidents where large data sets belonging to

well-known brands appeared after having been originally obtained years earlier. These incidents were frequently of a scale numbering in the millions, tens of millions or even hundreds of millions of customers. In some cases, the organisations involved were aware of a successful attack yet consciously elected not to disclose the incident. Many of the recent large breaches involved companies that *were* aware of unauthorised access to their systems, yet the scope of the intrusion was not known until years later when large volumes of data appeared in the public domain. In other cases, intrusions were entirely unknown until the organisation's data appeared publicly.

I've been personally involved in the disclosure of multiple incidents of this nature directly to the organisations involved. They're consistently shocked – *shocked* – that a breach had taken place and had not seen prior indicators that their data may have fallen into unauthorised hands. The passage of time frequently means that root cause analysis isn't feasible and indeed many of these systems have been fundamentally rearchitected since the original event.

It begs the questions – how much more data is out there? And what are we yet to see from events that have already occurred? We simply don't know nor is there any feasible way of measuring it. The only thing I can say with any certainty is that there is still a significant amount of data out there that we're yet to learn of.

A Perfect Storm of Data Exposure

Data breaches have been increasing in regularity and the incidents themselves have been increasing in terms of the volume of records impacted. There are a variety of factors contributing to what can only be described as a "perfect storm" of data exposure:

Firstly, as mentioned above, the rapid emergence of cloud services has enabled organisations and individuals alike to publish data publicly with unprecedented

ease, speed and cost efficiency. The low barrier to entry has meant that it's never been easier to collect and store huge volumes of information and very little technical expertise is required to do so.

Then we have the ever-increasing array of online services collecting data; social media sites, e-commerce, education, even cooking – every conceivable area of human interest has an expanding array of online services. In turn, these services request personal information in order to subscribe or comment or interact with others. As a result, the number of pools of user data on the internet grows dramatically and so too does the total attack surface of information.

The more recent emergence of the class of device we refer to as the "Internet of Things" or IoT is another factor. We're now seeing data breaches that expose information we simply never had in digital format until recently. In recent times, we've seen security vulnerabilities that have exposed data in cars, household appliances and even toys (both those targeted at children and those designed for consenting adults to use in the bedroom). All internet connected and all leaking data that didn't even exist in digital form a few years ago.

Data Maximisation as a Feature

Exacerbating both the prevalence and impact of data breaches is a prevailing attitude of "data maximisation", that is the practice of collecting and retaining as much data as possible. We constantly see this when signing up for services with requests for information that is entirely unnecessary for the function of the service itself. For example, requests for personal attributes such as date of birth and physical address, both data points that frequently provide no functional benefit to the service.

Further compounding the data maximisation problem is the fact that the retention period of the data usually extends well beyond the period in which the service is used by the owners of the data. (Indeed, even that term – "data

ownership" – can be interpreted to mean either the service retaining it or the individuals to whom the data relates.) For example, signing up to an online forum merely to comment on a post means the subscriber's personal data will usually prevail for the life of the service. There are many precedents of data breaches occurring on sites where those who've had their personal data exposed haven't used the service for many years.

Individuals' personal data is also frequently collected without their informed consent, that is it's obtained without them consciously opting in to the service and the purpose for which it's being used. Our data is aggregated, "enriched" and sold (often entirely legally) as a commodity; the people themselves have become the product and alarmingly, we're seeing the aggregation services themselves suffering data breaches both in the US and abroad. In this environment, it's the organisations holding personal data that control it, not the people to whom that data rightfully belongs.

I frequently hear from subscribers of HIBP that they have no recollection of using a service that's suffered a data breach. The alert they receive after the data is exposed is often the first they've heard of the service in many years. In fact, so much time has often passed that they frequently reject the notion that they were members of the site until they discover the welcome email in their archives or perform a password reset and logon to the service. The site was providing zero ongoing value to them yet it still retained their data and subsequently exposed it in a breach.

Data maximisation prevails as a practice for a variety of reasons. One is that it's increasingly cost effective to simply retain everything possible, once again due to the emergence of cloud services as well as rapidly declining storage costs. Another is that purging old data comes at a cost; this is a feature that has to be coded and supported. It also creates other challenges around technical constraints such as referential integrity; what happens to records such as comments on a forum when the creator of that comment has their record purged? Organisations view data on their customers as an asset, yet fail to

recognise that it may also become a liability.

Attempts by individuals to *reduce* their data footprint often lead to frustration. There's frequently no automated way of purging their own personal information and in some cases, organisations have even imposed a financial barrier in a "user-pays to delete" model. Even then, the purging of data from a live system is unlikely to purge that same data from backups that may stretch back years and we've seen many cases of the backups themselves being exposed in breaches.

We need to move beyond an attitude of data maximisation and instead embrace the mantra of "you cannot lose what you do not have".

There's a Lack of Accountability and a Propensity to Repeat Mistakes

Time and time again, we see serious data breaches that impact people's lives around the world and we ask "Is this the watershed moment?" "Is this the one where we start taking things more seriously?" Yet clearly, nothing fundamental has changed and we merely repeat the same discussion after the next major incident.

There's a lack of accountability across many of the organisations that suffer breaches as they're not held strictly liable for the consequences. Despite the near-daily headline news about major security incidents, there remain fundamental shortcomings in the security posture of most organisations. They trade off the cost of implementing security controls against the likelihood of a data breach occurring and inevitably, often decide that there's not a sufficient return on investment in further infosec spend. This attitude contributes to both the frequency and severity of serious security incidents and without greater accountability on behalf of the organisations involved, it's hard to see the status quo changing. There's not enough incentive to do things *right* and not enough

disincentive to do them wrong therefore the pattern repeats.

Data Breach Redistribution is Rampant

An important factor exacerbating the impact of data breaches is the prevalence with which the data is redistributed once exposed. Data breaches often spread well beyond the party that originally obtained it and the ease with which huge volumes of digital information can be replicated across the globe means that once it's exposed, it spreads rapidly.

There are multiple factors driving the spread of data that has been breached from a system. One is commercial incentives; data breaches are often placed for sale in marketplaces and forums where they may be sold many times over. The personal information contained within these breaches poses value to purchasers ranging from the ability to compromise other accounts of the victims' (frequently due to the prevalence of password reuse unlocking other unrelated services) to value contained within the accounts themselves (such as the ability to acquire goods at the victims' expense) through to outright identity theft (the accounts contain data attributes that help attackers impersonate the victim). In short, there is a return on investment for those who pay for data breaches therefore it has created a thriving marketplace.

More worrying though in terms of the spread of data breaches is the prevalence with which they're redistributed amongst individuals. Data breach trading is rampant and I often liken it to the sharing of baseball cards; two people have assets they'd like to exchange so they make a swap. However, unlike a physical commodity, the trading of data breaches replicates the asset as each party retains their original version, just like making a perfectly reproduced photocopy. Most of those involved in the redistribution of this data are either children or young adults, doing so as a hobby. Often, they'll explain it away as a curiosity; they wanted to see if any of their friends (or sometimes, enemies) were involved.

Other times they're experimenting with "hash cracking", the exercise of determining the original passwords when a system stores them as cryptographic hashes. They rarely believe there are any adverse consequences as a result of redistributing the data.

The exchange of data breaches is enormously prevalent. Sites hosting hundreds or even thousands of separate incidents are easily discoverable on the internet; there's often terabytes of data simply sitting there available for anyone to download. Forums dedicated to the discussion of data breaches frequently post links to new breaches or old data which may have finally surfaced. These are not hidden, dark web sites, these are easily discoverable mainstream websites.

Exposed Data is (Often) Immutable and (Usually) Irrevocable

Many of the data classes exposed in breaches are immutable, that is they cannot be changed. For example, people's names, their birth dates, security questions such as their mother's maiden name or even the IP address they were using at the time (which can be used to geographically locate them and potentially tie them to other exposed accounts). Other data attributes may be mutable albeit with a high degree of friction; an email address or a physical address, for example. They may both change over time but the effort of doing so is high and it's unlikely to happen merely because that data has been exposed in a breach.

Paradoxically, the data that is most easily changed is frequently the data people are most concerned about. Credit cards, for example, are often referenced in disclosure statements as not having been impacted by a breach yet a combination of fraud protection by banks and the ability to cancel and refund fraudulent transactions whilst issuing a new card means the real-world impact on card holders is frequently limited and short lived.

Exposed passwords are also easily changed and the impact of them falling into unauthorised hands can be minimal, albeit with one major caveat: The prevalence of password reuse means that the exposure of one system can result in the compromise of accounts on totally unrelated systems. But the password itself is readily changed and unlike immutable personal attributes, doing so immediately invalidates its usefulness.

Frequently, I'm asked how someone's data can be removed from the web; they're a victim of a data breach, now how do they retrieve that data and ensure it's no longer in unauthorised hands? In reality, that's a near impossible objective, exacerbated by the aforementioned redistribution of data breaches. Digital information replicates so quickly and is so difficult to trace once exposed, there's no putting the data breach genie back in the bottle.

The Emerging Prevalence of OSINT Data and the Power of Aggregation

Data available within the public domain is often referred to as "Open Source Intelligence" or OSINT data. OSINT data can be collated from a range of sources including social media, public forums, education facilities and even public government records to name but a few. It's data we either willingly expose ourselves or is made publicly available by design. Often, the owner of the data is not aware of its publicly available presence; they inadvertently published it publicly on a social media platform or had it put on public display without their knowledge by a workplace or school. In isolation, these data points may appear benign yet once aggregated from multiple sources they can expose a huge amount of valuable information about individuals.

Data aggregation – whether it be from OSINT sources alone or combined with data breaches – is enormously powerful as it can result in a very comprehensive personal profile being built. One system may leak an email address and a name

in the user interface, another has a data breach and exposes their home address then that's combined with an OSINT source that lists their profile photo and date of birth. Suddenly, many of the ingredients required to identify and indeed impersonate the individual are now readily available.

The Impact on Knowledge-Based Authentication

Knowledge-based authentication (KBA) is predicated on the assumption that an individual holds certain knowledge that can be used to prove their identity. It's assumed that this knowledge is either private or not broadly known thus if the individual can correctly relay it then, with a high degree of confidence, they can prove their identity. KBA is typically dependent on either static or dynamic "secrets" with the former being the immutable data attributes mentioned earlier (date of birth, mother's maiden name, etc.) and the latter being mutable such as a password.

The risks associated with static KBA have changed dramatically in an era of data breaches and an extensive array of OSINT sources. Further to that is the frequency and effectiveness of phishing attacks which provide nefarious parties with yet another avenue of obtaining personal data from unsuspecting victims. In years gone by, personal data attributes used for verification processes had very limited exposure. For example, one's date of birth or mother's maiden name would normally only be known within social circles which in the past, meant people you physically interacted with. A government issued ID was typically only provided to professional services that had limited exposure.

Now, however, the availability of static KBA data has fundamentally changed yet its use for identity verification prevails. The threat landscape has progressed much more rapidly than the authentication controls yet we're still regularly using the same static KBA approaches we did before the extensive array of

OSINT sources we have available today and before the age of the data breach.

Closing

Data breaches will continue to grow in both prevalence and size for the foreseeable future. The rate at which we willingly share personal data will also continue to grow, particularly with an increasing proportion of the population being "internet natives" who've not known a time where we *didn't* willingly share information online. Increasingly, the assumption has to be that everything we digitise may one day end up in unauthorised hands and the way we authenticate ourselves must adapt to be resilient to this.

Comments

So the question is, what are we going to do about it? I think it's high time that companies starts evaluating their security practices and operating system developers starting thinking secure by default in their engineering schemes. Code correctness and quality is another key, the future may hold hope with the rise of security-oriented programming languages like Rust. However languages like C can be just as effective if the all the libraries and standards are in sync and the code has been written properly. A good place to start would be with the open source community. OpenBSD is at the leading edge of futuristic security technology and efficient system's architecture. Another notably project is the MUSL libc project. Code correctness, consisted source code auditing, security by default, and rebuilding of modern software from new languages and tools is the right solution. I also personally think that new-school programmers could take lessons from old-school ones in optimization. The older cats had less powerful hardware to work with and were forced to write streamlined, efficient code. I think the less is more attitude is key. For a computer I think the OS install image shouldn't exceed 350-400 mbs. The smaller and more steamlined the base, the less attack

surface there is. However, windows and apple (linux too in some cases) just keep heaping on more bloat, more layers. So to sum up, good ENGINEERING practices and free(libre) software is the only solution!

None of that is going to save you from leaving a database backup lying around in a folder which is browsable and accessible on your website.

Nothing is going to protect us from stupid people. But that is no different than a company's financial chief leaving a briefcase with a million dollars on a park bench somewhere. The dudes gonna get fired for incompetence. Sysadmins who do this stuff get fired for incompetence and you hire better trained people. I truly believe we possess the technology to make things REALLY hard if not impossible to break into. Hackers depend on laziness, and bad system engineering. Everyone is scrambling now to take security more seriously. But it going to take re-inventing the wheel in alot of cases and rebuilding things from scratch. Not heaping on all these security layers onto a framework that's shit and a mess to begin with. However it's expensive to start over, it really is, I totally understand. That's why all these hospitals and governments were still using windows xp because doing massive system upgrades are really expensive. Therefore it was easy to hit them with a malware shit-storm. However, if these folks would stop putting their faith into proprietary technology which has failed us time & time again, they'd find that the free(libre) options have alot to offer in terms of protection.

That's great that the sysadmin gets fired. Except it wasn't the sysadmin that was responsible for the shitty security. It was the management that cut the IT budget down to nothing, leaving 0 minutes to think about security because IT has to do the jobs of 1000 people with 5.

Epilogue

To this day, my congressional testimony remains the single most noteworthy event in my career. It adorns my bio every time I need to look professional and regularly accompanies descriptions of me in the press. But more than that, it was just a momentous occasion for me personally to have gone from starting this little data breach pet project dealing with tonnes of illegally obtained data to being worried about HIBP's legal viability to somehow, ending up at US-freaking-Congress standing in front of lawmakers talking about the cyber. It still blows my mind.

There are so many aspects of this event that still stand out for me today. The genuineness, kindness and appreciation shown by everyone from staffers to the congressmen and congresswomen themselves, for example. I was just watching the recorded testimony again and it immediately struck me how, well, nice the lawmakers were, especially when it came to thanking me for travelling that distance. I also learned through this experience that for the most part, the folks you'll see in that video if you go and watch it again know very little about the actual topic being discussed and rely heavily on their staffers to do the background research and prepare the questions. That's totally understandable too; they're going from a hearing on data breaches at one moment to a hearing on the opioid epidemic the next. I had an opportunity to spend some decent time with the staffers in the leadup to travelling over, immediately before my testimony then over lunch afterwards and they were really switched on. Very clever people that understood my world exceptionally well and it was just an absolute pleasure to spend time with them.

I suspect this experience also really helped strengthen future relationships I'd establish. The following year I'd be working with both the UK and Australian Cyber Security Centres (the NCSC and ACSC) to make HIBP data available to governments (more on that later in the book). The year after that I was in New York at the NYPD Cyber Intelligence and Counterterrorism Conference I mentioned earlier in the epilogue to the ethics blog post, delivering a keynote

in front of 500-odd law enforcement officers from 3-letter acronyms departments around the world talking about how much illegally obtained data I'd processed! I'm confident the congressional testimony played a big part in getting my foot inside all sorts of doors over the years to come.

STREAMLINING DATA BREACH DISCLOSURES: A STEP-BY-STEP PROCESS

Of all the activities involved in loading data breaches into Have I Been Pwned, by far the most laborious and most infuriating is disclosure. It's also a blocking process; I'm not going to put a breach into HIBP unless I'm pretty confident the impacted organisation is already aware of it (per the example in the epilogue of how I verify breaches, I don't always get this right). Either that or I have to go to sufficient lengths to disclose that if the impacted organisation later turns around and gets cranky with me, they're not (reasonably) going to have a leg to stand on.

That last point is the heart of why I wrote this blog post, namely that there were enough occasions where I simply couldn't get any response at all that I needed to be very clear about what steps I'd take before going public anyway. I kept finding myself in this really difficult position, which is that I have a heap of data, I've got subscribers I've made a commitment to notifying if I find them breached, but I can't load the data and send emails until disclosure is complete. It's an awkward situation and yes, I could have expedited the whole thing by just loading the data anyway, but I didn't want to jeopardise my ability to run the service in the future because a company gets all legal on me for not disclosing the incident to them first! Very "rock and hard place" sort of stuff.

15 JANUARY 2018

don't know how many data breaches I'm sitting on that I'm yet to process. 100? 200? It's hard to tell because often I'm sent collections of multiple incidents in a single archive, often there's junk in there and often there's redundancy across those collections. All I really know is that there's hundreds of gigabytes spread across thousands of files. Sometimes - like in the case of the recent South Africa situation - I could be sitting on data for months that's actually very serious in nature and *needs* to be brought public awareness.

The biggest barrier by far to processing these is the effort involved in disclosure. I want to ensure that any incidents I load into <u>Have I Been Pwned</u> (HIBP) are first brought to the awareness of the organisations involved and whilst that may seem straight forward, it's often quite the opposite. There are notable exceptions (such as <u>the recent Disqus disclosure</u>), but more often than not, it's a laborious process of varying success. Because this is something I do over and over again, I want to streamline the process and more than that, I want to seek community input.

Tell me if I'm doing this right. This post documents how I intend to handle serious incidents with real consequences and frankly, I don't want to stuff it up.

What I'm going to do below is document the process I follow then apply it to 3 separate breach disclosures of different types. They'll each culminate with data being loaded into HIBP, subscribers receiving breach notifications and the incidents possibly landing up in the media. The consequences for the organisations involved can be serious; we've just seen <u>VTech fined \$650k</u> for an incident I was involved in disclosing a few years back. I don't take this lightly I want to ensure there's broad consensus on the way to handle these incidents. Let me explain my approach.

Defining Levels of Escalation

One of the biggest challenges with disclosure is getting a response from the

organisation involved in the first place. Following my recent Congressional testimony, I wrote a series on <u>fixing data breaches</u> and in part 3 I covered <u>the ease of disclosure</u>, specifically how organisations could make the process easier for folks such as myself to report security vulnerabilities or indeed previous breaches. In that post, I spoke about people giving up when it gets too hard:

Many well-intentioned people simply give up and don't report serious security incidents when the effort is too high or the risk is too great. That has to change.

That's often the case when it comes to reporting security flaws, but when you're already sitting on the data that someone has taken from an organisation's system, it can play out quite differently. In this situation, the question is not whether disclosure happens or not, but rather how much effort I should go to before the data ends up in HIBP.

Most of this post is going to talk about levels of escalation, that is the phases I go through with each one increasing in effort. I'm going to start with the most obvious channels for reporting a breach and move through increasingly indirect ways get the organisation's attention. These channels and the timeframes in which I escalate through stages are essential - please comment on them!

Let me detail those levels then jump immediately into the 3 use cases I'll be applying the process to. Do note that I'm writing this *before* loading those 3 breaches; that'll happen immediately after I publish this blog post and they'll each reference it.

Level 1: Published Security Contact Info

Ideally, organisations acknowledge that security incidents may occur and they provide dedicated channels through which to communicate with them. They may adhere to a convention or be published elsewhere on their website:

- 1.Look for a security.txt file at /.well-known/security.txt (read more about this)
- 2.Look for a security vulnerability reporting policy (Tesla has a great example)
- 3.Look for a published security email address

The reality is that it will be a rarity to find this information, but it's always the first preference and I wanted to include it here as encouragement to organisations. If no published security contact info can be found, immediately fall through to level 2:

Level 2: Published General Contact Channels

General contact channels are less ideal as organisations receive all sorts of communication via them. Reports of serious security incidents are likely to land next to Viagra spam; there's a poor signal-to-noise ratio and it may mean an expeditious response will not be forthcoming.

- 1.Look for a contact us form
- 2.Look for publicised email addresses
- 3.Look for social media accounts that accept private communications
- 4.Email "standard" addresses: security@, admin@, webmaster@, support@, postmaster@, hostmaster@

All the above channels may be used simultaneously - the same message can be sent to each. A period of 3 business days should be allowed for a response before proceeding to the next level.

Edit: Thanks to the comment from Stuart, point 4 was added after his feedback.

Level 3: Indirect Contact Information (Optional)

Once level 3 is reached, "reasonable attempts" have been made to contact the company. By this time, either no readily accessible contact information has been found or the company has not been responsive to the report. Level 3 involves either increasingly hard work or a decreasing likelihood of success so is flagged as optional.

- 1.WHOIS records
- 2. Company employees via LinkedIn
- 3.Other sources of OSINT contact info
- 4.Email addresses on the domain already in HIBP

If level 3 is used, provide a period of 3 business days for a response before proceeding to the next level.

Edit: Thanks to the comment from Rob, point 4 was added after his feedback.

Level 4: Public Requests for Information

A public request for a security contact can be construed as implying the company has a security issue. Whilst this may be true, alerting the general public to this (even without providing details of the issue) before the company themselves should be a last resort.

1.Request security contact via social media

By this time, 3 to 6 business days have already passed (we're potentially more than a week in) and reasonable efforts have been made to contact the organisation. Public requests via the likes of Twitter get many eyes very quickly so if there are no solid leads within 24 hours, progress to the final level.

Level 5: Full Public Disclosure

By this time, there has usually been many hours invested (including verifying the incident), security contacts have been sought, all publicised contacts tried and sufficient time for a response has passed. The organisation simply won't respond and it's now time to publicly disclose the incident via HIBP.

Disclosure in Practice 1: The Fly on the Wall

It all started with this tweet:

Just hijacked some big MySQL database server containing 53K credit card details with complete CVV2 happy new years to the 4 million users pic.twitter.com/pXda5DbNCz

-- *Taylor* (@0x55*Taylor*) *December* 31, 2017

The data was sent to me and after inspecting it, I found identified 84k email addresses in the breach. The data included passwords stored in plain text and a quick password reset check on a Mailinator account delivers the precise password in the breach to the public mailbox. It's almost certainly legit, let's

move onto the disclosure process.

Level 1: No security.txt file at <u>theflyonthewall.com/.well-known/security.txt</u> nor a security vulnerability reporting policy or a security email address anywhere.

Level 2: Generic contact form located and completed, submitted at 23:44 UTC on Jan 6 with the following message:

Hi, my name is Troy Hunt, I'm a security researcher you can read about here: https://www.troyhunt.com/about/

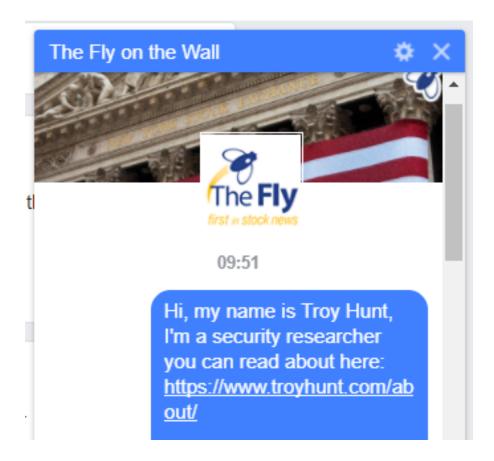
I was recently sent a data breach alleged to have come from theflyonthewall.com and upon verifying it, I believe it's legitimate. The data indicates that your website has had a large amount of data extracted from it, including credit card numbers. Could someone in a security capacity please get in touch with me via email so I can provide further information.

Confirmation:

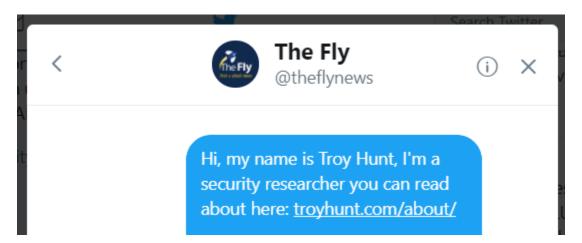
Thank you for you submission!

The Fly makes every effort to respond to all emails. If you do not hear back from us within one business day, please email us again.

7 minutes later, I also sent the same message via their Facebook page:



And via DM on their Twitter account:



Level 4: I've given it more than 5 days (well beyond the 3-day target) and there's been zero response. Remember, I've tried the contact form, Facebook and Twitter by now so it's not from lack of trying. At 08:58 UTC on 12 Jan, I send out a tweet:





Anyone got a security contact at @theflynews? I've been trying to get in touch with them about a serious security incident for 5 days now and can't get a response by any published channel.

♥ 8 10:58 PM - Jan 11, 2018

0

Level 5: No replies of substance or any other contact made, the data is loaded into HIBP at 06:47 UTC on 15 Jan and 386 individual and 815 domain subscribers are notified of the incident.





New breach: Stock market news website The Fly on the Wall had 84k unique email addresses breached along with purchase histories and credit cards. 64% were already in @haveibeenpwned. Read more: troyhunt.com/streamlining-d...

○ 57 8:47 PM - Jan 14, 2018

0

Disclosure in Practice 2: Open CS:GO (Counter-Strike: Global Offensive)

I'm handed a 10GB MySQL backup file with 512k unique email addresses titled csgo_20171128.sql which allegedly came from opencsgo.com. That URL promptly redirects to dropgun.com but the site doesn't exhibit any of the usual

verification vectors I'd use because access is only granted via external services (Facebook, Twitter, Steam login). I fire off an email to the 30 most recent HIBP subscribers I find in the data (there are 816 subscribers in total in the data):

Hi, I'm emailing you as someone who has recently subscribed to the service I run, "Have I been pwned?"

Your email address has appeared in a new data breach I've been handed and I'm after your support to help verify whether is legitimate or not. I'd like to be confident it's not a fake before I load the data and people such as yourself receive notifications.

If you're willing to assist, I'll send you further information on the incident and include a small snippet of your (allegedly) breached record, enough for you to verify if it's accurate. Is this something you're willing to help with?

For verification, I'm on the about page of the site: https://haveibeenpwned.com/About

Within 14 minutes, the first person has responded, I've sent them information on their account and they've confirmed the accuracy. Others chime in over the next 24 hours, each confirming the legitimacy of their data and that they had indeed used the service.

So let's apply the process again:

Level 1: No security.txt file at <u>dropgun.com/.well-known/security.txt</u> nor a security vulnerability reporting policy or a security email address anywhere.

Level 2: Onto level 2, I prepare a message for them:

Hi, my name is Troy Hunt, I'm a security researcher you can read about here: https://www.troyhunt.com/about/

I was recently sent a data breach alleged to have come from opencsgo.com (which now redirects to dropgun.com) and upon verifying it, I believe it's legitimate. The data indicates that your website has had a large amount of data extracted from it, including personal information. Could someone in a security capacity please get in touch with me via email so I can provide further information.

This goes out via both email to <u>their published address</u> and <u>Twitter DM</u> at 21:04 UTC on Jan 7:



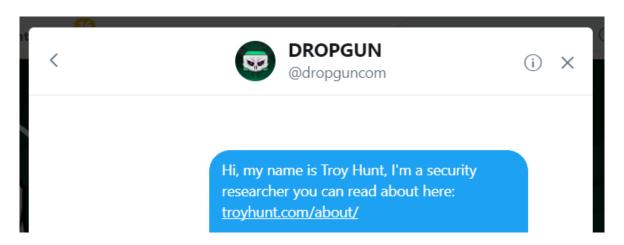
Mon 8/01/2018 07:03

Troy Hunt

Security Incident Impacting Open CS:GO / Dropgun

To 'support@dropgun.com'

Hi, my name is Troy Hunt, I'm a security researcher you can read about here: https://www.troyhunt.com/about/



Their Facebook page doesn't allow messages.

Level 4: Same deal as The Fly on the Wall and at the same time as tweeting out for public support re them (08:58 UTC on 12 Jan), I send out a public tweet about Open CS:GO:



And again, nada.

Level 5: Just like The Fly on the Wall, there's no replies of substance or any other contact made so the data is loaded into HIBP at 06:17 UTC on 15 Jan and 826 individual and 243 domain subscribers are notified of the incident.



Disclosure in Practice 3: Lyrics Mania

I receive a 4MB CSV file with 109k lines of email addresses, usernames and plain text passwords. The Lyrics Mania website is a basic site listing song lyrics and has no obvious means of registration or authentication so again, that makes verification hard. A bit of Googling locates a login page on another subdomain (oddly, without HTTPS when the primary site has it), but a password reset with a Mailinator address returns the same page (seems to just reload it and not properly submit). I literally diff the responses and there's nothing there to confirm or deny the request has even been received and no email lands in the Mailinator inbox.

It's borderline in the "too hard basket" for the return given the (relatively speaking) small number of accounts, but the plain text passwords make it a more serious incident so here I go again. I find 233 accounts of HIBP subscribers

and email the most recent 30 as I'd just done with the CSGO breach. I immediately get back 2 "out of office" replies, both in German. Over the next day, multiple other parties reply and confirm their data is legitimate and that they've used the service. It's real, so let's apply the process again.

Level 1: No secuity.txt file at <u>lyricsmania.com/.well-known/security.txt</u> nor a security vulnerability reporting policy or a security email address anywhere.

Level 2: Similar message to before:

Hi, my name is Troy Hunt, I'm a security researcher you can read about here: https://www.troyhunt.com/about/

I was recently sent a data breach alleged to have come from Lyrics Mania and upon verifying it, I believe it's legitimate. The data indicates that your website has had a large amount of data extracted from it, including personal information. Could someone in a security capacity please get in touch with me via email so I can provide further information.

I find a contact page, fill out the form and submit it. It doesn't go so well:



Oops, there is an error Page not found

You can search Lyrics by using the search bar above, or visit the homepage



There's <u>a Facebook page</u> but that doesn't allow messages to be sent. Without any means of contact directly published via their website, contacting these guys is becoming increasingly hard. Let's fall through to level 3:

Level 3:

<u>Their WHOIS record</u> has privacy enabled so there's no contact info of any use there. I search through LinkedIn and find <u>one person with "Lyrics Mania" on their profile</u>:



Midhun Vincent • 3rd Senior Android Developer at Eight Signs FZ LLC

Eight Signs • Sathyabama University
United Arab Emirates • 500+ &

Send InMail

It looks like it's a hobby project for the guy:

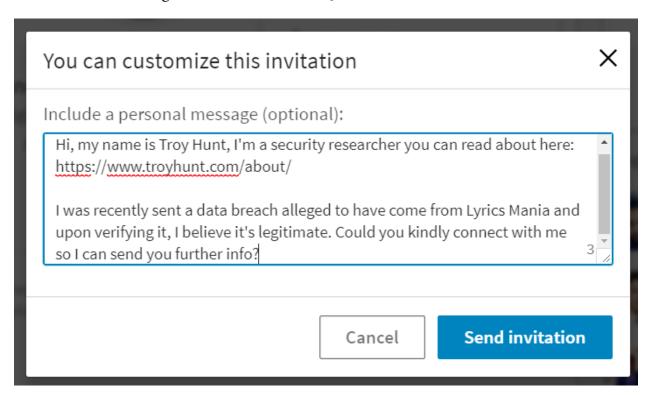
7 Projects

Lyrics Mania Nov 2014 – Present

Lyrics Mania is a mobile app with the biggest lyrics database on the market. Lyrics Mania lets you search for lyrics in a wide catalog. With the integrated music player, you can listen to your music and get lyrics in real time while you're enjoying your favorite songs.

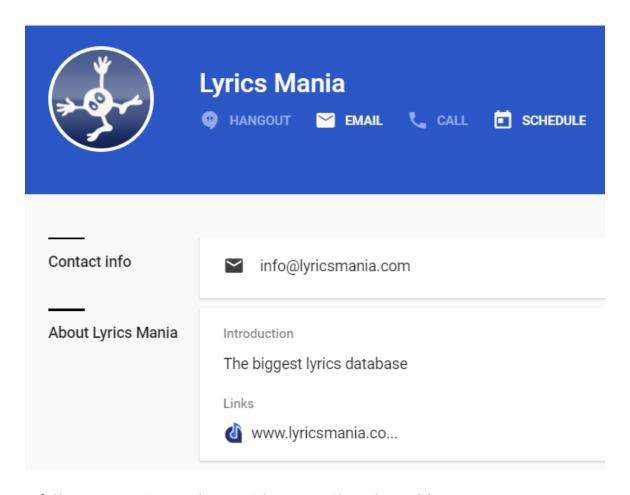
Obviously, I'm going to be sympathetic to someone who's stood this up in their spare time for fun, but by the same token we're looking at over 100k passwords that people are using on their email accounts, banking and who knows what else. It's a serious incident and it has to be seen through.

I send him a message at 21:37 UTC on 7 Jan:



✓ Your invitation to Midhun Vincent was sent! View Profile

I also find that the Lyrics Mania Google+ page has an email address:



So I follow up 4 minutes later with an email to that address:



Mon 8/01/2018 07:41 Troy Hunt

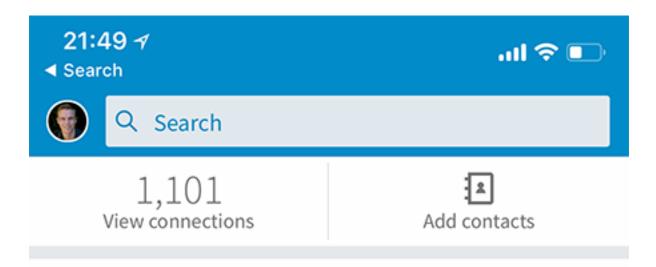
Security Incident at Lyrics Mania

To 'info@lyricsmania.com'

Hi, my name is Troy Hunt, I'm a security researcher you can read about here: https://www.troyhunt.com/about

I'm not overly confident of an expeditious response, but I'll give it 3 business days and see what happens.

About a day and a half later, Midhun connects:



Invitations (131 new)



Midhun Vincent accepted your invitation

It's late for me so first thing the next morning my time (20:16 UTC on 9 Jan), I message him the following:

Hi Midhun, thanks for connecting.

I run the data breach notification service "Have I Been Pwned" (HIBP): https://haveibeenpwned.com/

People regularly send me data exposed in data breaches and someone recently sent me 109,219 records of email addresses and plain text passwords allegedly from Lyrics Mania. A number of HIBP subscribers were in there and I've confirmed with them that the data is indeed legitimate. For your verification, here are 3 of the Mailinator addresses in the file:

[redacted]@mailinator.com

[redacted]@mailinator.net

[redacted]@mailinator.com>

Is there somewhere I can send you the data so that you can review it? If it is indeed legitimate, you'll need to get in touch with those in the

breach and notify them, especially given there's plain text passwords which will be usable on other services too.

Less than 48 hours to get acknowledgement isn't too bad, it's the fastest turnaround of all 3 and it's the only one that's a hobby project too. Neither of the other commercial operations have responded yet.

Level 5: Unfortunately, more than 5 days later and even after accepting my LinkedIn invitation (which clearly flagged the security incident), I've had nothing back from Midhun. I load the data into HIBP at 06:34 UTC on 15 Jan and 233 individual and 284 domain subscribers are notified of the incident.



New breach: Song Lyrics website Lyrics Mania had 109k records breached including usernames, email addresses and plain text passwords. 78% were already in

@haveibeenpwned. Read more: troyhunt.com/streamlining-d...

Incidentally, I know some people may be critical of my naming Midhun here. But if I'm to be transparent about communications then it's going to result in identifying those I've attempted to contact and who are responsible for the data. And it *is* a big responsibility too because as I said earlier, these are plain text passwords for all sorts of important accounts people will have reused them across and accountability must be taken when an incident like this occurs.

Comments

In previous posts about breach disclosure you've mentioned journalists (and alike) as a

means of getting in touch with the right people and getting them to respond. Could that be something to include in above list?

Then a slightly offtopic comment and a random idea that popped into my mind while reading, it might be a stupid idea, but nonetheless here it is:

You've written about breach verification multiple times and multiple times you mentioned asking the most recent HIBP subscribers to help and verify their data. If I where ever in a position where I could help with verification, I'd be happy to help and respond immediately, but I'm by no means "recent". Maybe an idea where people could mark themselves (opt in obviously) as someone who's willing to help which means you'll have a pool of people who will most likely respond to such requests. Idk, might be too much effort to setup such a system, just a random thought I had...

Troy: I actually did discuss The Fly on the Wall with a journo. He didn't run with the story because frankly, breaches like these are now old news. I'll definitely still work with journos, but I'm finding they need increasingly big stories to even justify the effort.

As for subscribers willing to help, I find that grabbing the 30 most recent works well; HIBP is something they remember and I always get enough responses to verify what I need to.

A few comments:

- Think about the effectiveness of LinkedIn based on your direct experience of lack of follow up, and how dev type people often see it more as a source of recruiter spam rather than a source of actual worthwhile contact.
- Consider searching the hacked email addresses for emails in the domain that got breached and finding obvious employee accounts. Correlate with LinkedIn if necessary.

• Consider writing a easily readable (i.e. this post as a TL;DR doc) disclosure statement and FAQs targeted at sites that are been pwned rather than users. Points like: how certain are you that this is a problem, why is this a problem for the site, what should they do to confirm, what will you do if you don't hear from them, timelines etc.

You should definitely say what you will be doing - i.e. very public disclosure in about a week from now whether or not they reply.

I would also include some public links by major news orgs about HaveIBeenPwned for credibility.

Troy: I deliberately don't mention HIBP in the communication as that's not necessarily what will be the outcome (i.e. the Red Cross Blood Service didn't go in). I'm also conscious that it could be very confronting if I did; "Hey, I have your data and I'm going to load it into a publicly facing system". I'm trying to approach this with a little more tact than that. although I do understand the point you're making.

I think your mail looks a bit spammy because it has an URL in it advertising your own web site. I understand that you want to provide more information about you, but still ... unfortunately, I don't have a good improvement suggestion.

Also, I think many companies need some time to process the information internally (mail receiver contacts the securtiy guy, who is out of office, who responds two days later, then they have to contact their very busy manager, as it is a huge leak he has to contact his manager, etc. pp. you get the picture). I think it can easily take about two weeks to internally process such a mail and come up with an answer.

Troy: There's a few good points here so let me address each one individually:

1.The URL is important context: that "about" page explains who I am and

what I do and whilst I understand your point, like you I don't have a better suggestion!

- 2.I haven't documented it here, but if I received an OOO then that may impact the timelines (i.e. wait until they return or contact an alternate person mentioned in the OOO)
- 3.Even with delays due to internal processing, there's no good reason not for someone to reply promptly to me once receiving the message

On that last point, I've had cases where internal processing has taken *way* too long. Nissan is a good example where I got a prompt response via email and had good initial dialog, but their failure to act one month later led to this: https://www.troyhunt.com/co...

Epilogue

Half a year after writing this blog post I had a great test of the process courtesy of Adult-FanFiction. Never heard of them? Imagine you really liked reading the erotic adventures of a vampire and his Ewok sex slave. Or something to that effect, it was all written erotica about fantasy characters and ok, that's not really my thing but hey, who am I to judge. A data breach is still a data breach no matter who it happens to, and I wanted to make sure it was disclosed, loaded and HIBP subscribers notified.

I checked back through my emails and it was the 19th of July 2018 that I contacted them at both their technicalsupport@ and forums@ email addresses, both of which were published publicly. I included the usernames and password hashes from the first 3 rows of the breach, more than enough data to verify that they had indeed been breached. I followed the process I'd outlined in this blog post and failing to get any response from them whatsoever, published the data

on the 6th of August. This was much longer than the timeline I'd outlined earlier, but it was a small breach and I'm a busy guy. So, breach loaded, job done, onto the next thing. Until...

Someone posted about it on their forum after receiving an HIBP notification email. Someone else received one too and also chimed in. And then the mods joined the discussion and although this thread was later deleted, I took some handy backups I'll quote from here:

"No, it's not true. One would think that instead of scaring users, these companies would contact the site itself" - DemonGodess (Head Evil Tech Wench)

For fuck's sake. You mean contact them like via the email addresses they published?! Plus, just use your brain for a moment: you've got multiple members of your service who advised they'd received a breach notice regarding your service, what are the odds that would happen by accident? And then, finally, I had a response to my earlier email land in my inbox:

"None of this data exists in my database. Sorry for the late response, I've been working hellacious hours at my day job"

She followed up 2 minutes later with a second email:

"Now that I've verified this doesn't exist in my database, I expect you to remove the site from your data breach website. Now, please."

No "pretty please"? Regardless, I get back to her with a much more comprehensive email including a screen cap that was sent to me by the person who reported it showing a SQL injection attack in the URL. I also sent back a list of the table names that had been extracted, again by the person who sent me the data. In the meantime, another mod on the forum is doubling down on the stupid:

"There is no gain in hacking us. We're not an e-commerce site, but if you are looking to terrify people into buying security software, this is an extremely effective gambit" - BronxWench (Mighty Dragon Wench)"

You know what makes people want to hack you? You're on the internet. That is all. That's the starting point. Beyond there, some sites are more attractive to break into than others, but that's your minimum criteria and it's one they met. She continued:

"You're right, in the he sells private workshops and courses. During those, you will be offered the opportunity to purchase security packages to prevent data breached from impacting you."

She got the first sentence right and went downhill from there. I'm relaying this in detail here because it amply describes the sort of rubbish I have to regularly deal with in order to run this service. One final comment from BronxWench:

"Perhaps because we don't have high priced lawyers who will take him to court, or perhaps he thinks our members won't report him for publishing their data without permission? That's what I intend to do."

And that was the last I ever heard from either of them. I never felt genuinely threatened, but this experience perfectly illustrated why I wrote this blog post and why I feel the need to approach disclosure with the diligence I do. One day it might not be an Evil Tech Wench or a Mighty Dragon Wench writing Chewbacca porn that wants to take legal action, it might be a high-powered law firm representing a Fortune 500 company and that, I suspect, would be a very different experience.

THE UK AND AUSTRALIAN GOVERNMENTS ARE NOW MONITORING THEIR .GOV DOMAINS ON HAVE I BEEN PWNED

In June 2017 I was in London doing a talk at the .NET user group at Skills Matter. It was massive for a user group - there must have been 200 people there - and I was doing my usual infosec show, this one titled "Clouds, Codes and Cybers". I did the talk and then went to Q&A with the usual collection of questions from the audience. And then someone asked one about governments which effectively boiled down to "isn't the government just always trying to get into all our stuff and screw us over"? The event wasn't recorded so I can't quote myself word for word, but it boiled down to this:

"Governments have an extremely difficult job to do in balancing the privacy we all enjoy with the freedoms we all enjoy. Their job is getting harder and harder as we have more ubiquitous encryption, and their actions are scrutinised more than ever (the Snowden leaks were still very fresh in everyone's minds). When we learn that the gov has overstepped boundaries there's guite rightly outrage and its front-page news. However, when they do a great job, we rarely hear about it because nothing happens when they do their job well! So understandably, our perception is shaped by the narratives we all observe. But what we don't see is the dedication of so many really top-notch people working in gov cyber departments. Genuinely nice people you'd want to go and have a beer with, people that are paid a fraction of what they'd earn in private enterprise, but they do what they do because they want to make a positive difference in the world. In my experience and without exception, everyone in a gov cyber department I've ever met anywhere I've travelled around the world has been a standout individual and it's sad that perceptions like the one expressed in this question prevail."

I wrapped up the Q&A, grabbed a beer and wandered around. A bloke came up to me and we started chatting; he was from the National Cyber Security Centre (NCSC), a department of the Government Communications Headquarters (GCHQ) which is sometimes referred to as "Britain's Spy Agency". He was really appreciative of the response I'd given, and I got the distinct impression he was more used to hearing people deride the role he played in protecting the public than he was hearing people being ingratiating. He was also one of those genuinely nice gov people I'd previously enjoyed having beer with so over a couple of cold ones, we got talking. That's where the idea to open up HIBP to govs was born and 9 months later, I published this blog post.

02 MARCH 2018

If I'm honest, I'm constantly surprised by the extent of how far <u>Have I Been Pwned</u> (HIBP) is reaching these days. This is a little project I started whilst killing time in a hotel room in late 2013 after thinking "I wonder if people actually know where their data has been exposed?" I built it in part to help people answer that question and in part because my inner geek wanted to build an interesting project on Microsoft's Azure. I ran it on a coffee budget (the goal was to keep the operating costs under what a couple of cups from a cafe each day would cost) and I made it freely accessible. And then it took off.

As this service has grown, it's become an endless source of material from which I've drawn upon for conference talks, training and indeed many of my blog posts. I've written extensively about how HIBP has grown over the years and doing so has been a cornerstone of the philosophy of how I've run the service - with maximum transparency. My view has always been that it's in everyone's best interests to be crystal clear about how I run this, especially when you consider the circumstances of how most of this data was leaked in the first place. And this is precisely why I'm writing this piece - to talk about how I'm assisting the UK and Australian governments with access to data about their own domains.

Just to scroll back for a bit of context, anyone who owns a domain can do a free domain search on HIBP. There's a verification process where control of the domain needs to be demonstrated (email to a WHOIS address, DNS entry or a file or meta tag on the site), after which all aliases on the domain and the breaches they've appeared in is returned. At the time of writing, over 110k domain searches have been performed *and verified*. These searches span every imaginable class of domain; financial institutions, aerospace, healthcare, adult entertainment and based on a very rough check just now, more than a quarter of all Fortune 500 companies as well. Amongst those verified domain searches are government departments and they too are enormously varied; local councils, legal and health services, telecoms and infrastructure etc. The thing is, loads of government departments within different countries have all been running these searches independently and that means an awful lot of duplication of effort has been going on. This post talks about how I'm addressing that.

Over recent times, I've had a bunch of opportunities to talk to folks in various government roles. My congressional testimony in the US was a very public example of that, less so are the dozens of conversations I've had in all sorts of settings including during conferences, workshops and over coffees and beers. The subject outlined above (loads of government departments independently using HIBP) came up in a number of those meetings, so we decided to do something about it. Not only did we want to consolidate all those existing independent departments doing their own thing, we wanted to expand the scope to all government departments. So, this is what we've done:

As of now, all UK government domains are enabled for centralised monitoring by the <u>National Cyber Security Centre</u> (NCSC) and all Australian government domains by the Australian Cyber Security Centre (ACSC).

The way we're doing this is by using the commercial model within HIBP and scoping it to a limited set of whitelisted domains. Whilst I've written and spoken publicly about commercial services in the past, I've avoided promoting them per se which is why you won't find anything on the HIBP website or any up-sells in

emails or anything along those lines. I've always wanted HIBP to be first and foremost a freely available service for email and verified domain searches and particularly in this industry, it's very easy for financial motives to taint the ethics of how this data is dealt with. To that point, I've made this available to the NCSC and the ACSC without any commercialisation whatsoever - they get it for free. There are many reasons why that made sense to do, one of which is that it unifies a bunch of existing free searches that I mentioned above. Another is that frankly, we really *want* governments to do their best to protect the folks working in their departments; many of them are working in capacities that help protect our respective nations from all sorts of threats and increasingly, as we all know, that means online threats as well.

Getting back to the mechanics of things, the respective govs are using the commercial HIBP model in a tightly scoped fashion. For example, the UK government can query any .gov.uk domain on demand and the Aus government can query any .gov.au domain on demand. They can both also query a small handful of whitelisted domains on different TLDs, for example, The Commonwealth Scientific and Industrial Research Organisation (CSIRO) runs on csiro.au so that domain is whitelisted for the ACSC in addition to the .gov.au TLD. What this means - and this is enormously important - is that the NCSC and ACSC can't turn around and query, say, troyhunt.com. The only access they have is to domains that their people working in those departments could query anyway via the existing free domain search model, we're just consolidating it all into a unified service.

As part of that service, they'll also be using the existing notification service that commercial subscribers have access to. This is a webhook model which calls back into an endpoint the respective governments host. Every time an alias on one of their domains is seen in a new data breach or a paste, the incident is automatically posted to them. It means that within minutes of one of their email addresses being found and loaded into HIBP, they'll know about it. That's really important in terms of giving them the ability to respond quickly and by unifying all those existing one-off domain searches, the respective governments will be

able to immediately see when an incident has a potentially broad impact. This can be especially important when you consider data breaches such as Dropbox; many organisations of all kinds suddenly learned that a bunch of their people had cloud storage accounts under their corporate email addresses so you can imagine some of the discussions that subsequently ensued.

So that's what's been set up in HIBP for the UK and Aussie governments. They both respected my desire for transparency and understood why it was important for me to write about it publicly. I'm happy that this effort continues the philosophy I've stuck to since the early days of HIBP - that the service should help people do good things after bad incidents occur and that it does so as transparently as possible

Lastly, I want to touch on something a bit tangential and that's to point out some of the great work these agencies are doing to try and improve online life for the rest of us. I've regularly quoted the NCSC in particular, for example there's a bunch of their work in my recent blog post about authentication guidance for the modern era. I love that we have a government department making recommendations such as "only ask users to change their passwords on indication of suspicion of compromise" because it validates what an increasingly large number of us in the security industry have been saying for so long. The NCSC piece on let them paste passwords is another favourite; the blog post of mine they reference there has led to a regular cadence of people pointing out sites that don't adhere to this guidance and having a government resource for me to point offending companies at is enormously valuable. Likewise, in Australia we have <u>Stay Smart Online</u> which provides a bunch of consumer-level information about precisely what the name of the site suggests. They regularly highlight emerging scams and other digital threats to everyday Aussies as well as creating practical guidance for the fundamentals such as guidelines for <u>creating strong passwords</u> (I particularly like the fact this draws on the NIST recommendations I included in the aforementioned authentication guidance blog post). I'm very happy that HIBP is now a resource the UK and Aus governments can draw on to help their people help all of us live happier (and

hopefully less pwned!) online lives.

Epilogue

Since the UK and Aussies governments began using HIBP, dozens more have followed. A bunch in Europe, the US and Canada, South America, the Middle East and in many of those cases, they're via people I spent time with face to face. I visited and toured around their facilities, chatted about all manner of cyber things and often spent time grabbing dinner and drinks. In several cases I'd regularly catch up with these folks when visiting their respective countries, not because of anything to do with HIBP, but rather because they're just super interesting people doing really important work. I feel more strongly about the answer I gave in 2017 than ever.

The gov access to HIBP is a minor overhead for me. Beyond the initial setup, I spend an amount of time each month supporting the relationships which rarely exceeds an hour. But it means a lot to them, and I've heard multiple stories of how data from these relationships has made a positive impact on everything from education to actively combatting account takeover attacks. And it's been good for HIBP too; being able to publicly talk about govs leveraging the data I've accumulated in the service has been invaluable, especially in the wake of law enforcement action against the likes of LeakedSource. It's honestly been a real buzz for me personally to see a bunch of these governments excitedly announce the relationship on their social media with the word "pwned" proudly appearing alongside official government logos

Perhaps this is a service I could have charged money for. Maybe I could have done as so many other organisations have done and turned their cyber things into a revenue stream funded by public coffers. But what I love about HIBP is that I've never needed to do that so I've just done what feels right rather than what might be in my financial best interest. I love that this philosophy has defined how I run the service and IMHO, it's all the better for it.

THE LEGITIMIZATION OF HAVE I BEEN PWNFD

You may have noticed a bit of a recurring theme across some of the posts of the last 6 months I've selected for inclusion in this book: HIBP was becoming a more mainstream, well-known service. The ethics blog post, the one on Congress and now this one were all intended to help position HIBP in a positive light. Especially in the context of some of the shadier things I was seeing (the Leaked Source takedown was just over a year earlier), I felt this constant need to justify HIBP's role in doing good things, despite the amount of bad things that were being done with precisely the same data sets.

21 MARCH 2018

here's no way to sugar-coat this: <u>Have I Been Pwned</u> (HIBP) only exists due to a whole bunch of highly illegal activity that has harmed many individuals and organisations alike. That harm extends all the way from those in data breaches feeling a sense of personal violation (that's certainly how I feel when I see my personal information exposed), all the way through to people literally killing themselves (there are many documented examples of this in the wake of the Ashley Madison breach). Plus, of course, there's the ginormous financial impact; <u>TalkTalk claims their 2015 hack cost them £42M</u> and I've heard first-hand from those inside other companies that have suffered data breaches about just how costly they've been ("many millions of dollars" is very common).

Since day 1 of running HIBP, I've been overtly conscious of the shadiness of the realm within which it operates and consequently, I've done everything I possibly

could to position it in the most ethical light possible. Transparency has been a huge part of that effort and I've always written and spoken candidly about my thought processes, how I handle data and very often, the mechanics of how I've built the service (have a scroll through the HIBP tag on this blog for many examples of each). Indeed, my own comfort level with the legitimacy of running this service has changed over time and that's really what I wanted to talk about here in this post: where it's come from, where it is today and how over time, it's been increasingly legitimised. This has changed most fundamentally in the last year and a bit so let me start there.

The Industry Cleaned Up a Lot in 2017

I very consciously avoided talking about it publicly at the time (largely because I didn't want to draw attention to it), but particularly around late 2016 and very early 2017, I was quite concerned with the broader genre that is data breach search services. The reason for this was that there were an awful lot of them operating in a very shady space attracting the wrong sorts of attention. Chief among these was <u>LeakedSource which was eventually taken down in Jan last year</u>.

Let's just recap on the value proposition of this service for a moment: for as little as 76c a day, you could subscribe to LeakedSource and view the raw data from a breach. What this meant was that people could pay their cash and access the personal data of *anybody*, including myself (sent to me by someone who bought access):

Filtered Results from DbForums.com
Hacked on: 2016-07-04
Encryption method for passwords: Vbulletin
username: troyhunt
hash: c7452c5009737e8b9cac
email: troyhunt@hotmail.com
register_date: 1359347497
last_login: 1359347805
birthday: Market Ma
ipaddress:
salt: xgs`9t }5*"3pa3?]xzn:h'WoN]\[h

No verification of ownership, no censoring of results just the full (often sensitive) personal information of victims of data breaches. The operator justified the service by saying that the data was "all freely available online" and that in order to search for someone else's data, you must "have his written permission you may do so". Yet at the same time, they had absolutely no checks on data ownership and actively advertised the service on hacking forums which, of course, attracted precisely the sorts of customers you'd expect:

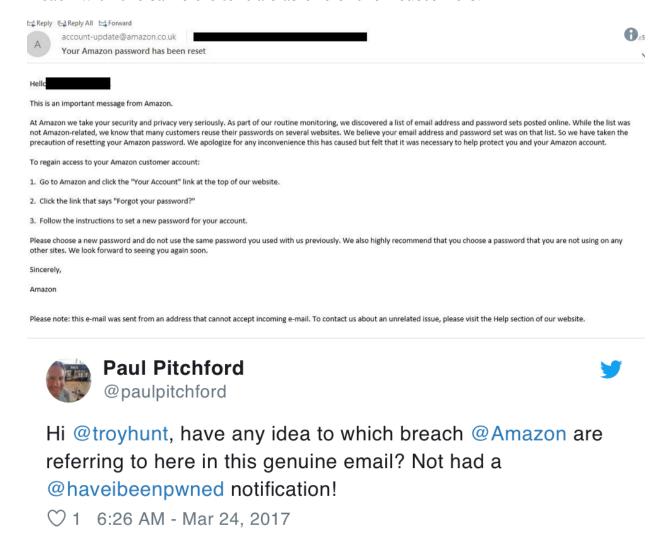


I wish leakedsource had an app so I can jack cod channels on the go!

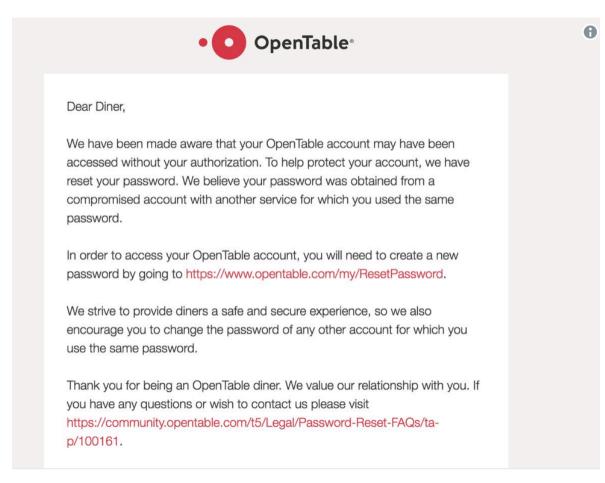
In Jan last year, that all came to a screeching halt. It took another year after that before the Canadian Mounties charged an Ontario man with a host of offences including "trafficking in identity information". (The video with RCMP Staff Sgt. Maurizio Rosa in it is worth a look, note the comment regarding the purchasing of data too, a practice I've always been vehemently against due to the incentives it provides to hackers.) He's yet to face court and answer those charges, but it doesn't look like it's going to work out real well for him.

I don't want to just focus on LeakedSource though, whilst it was the most notable at the time there were many others operating in a similar space (<u>Leakbase was another that "went dark" in 2017</u>). Yet at the same time, other organisations increasingly began using breach data *to do good things* and this is where my comfort level really started to shift.

A good example of this is the notifications Amazon sends when they find a data breach with the same credentials as one of their customers:



I love this because it's proactive: Amazon have grabbed data that's circulating and taken proactive steps to protect both their customers and themselves. OpenTable also seems to be pulling data from other breaches:







.@troyhunt @GossiTheDog did @OpenTable get popped?

O 12:31 PM - Mar 7, 2018

(Side note: getting the wording of these emails right is absolutely critical, as is evidenced by the accompanying tweet which casts suspicion over OpenTable's security posture.)

LinkedIn also does the same thing, this one sent to me by a follower:

Linked in

Hi√ ,

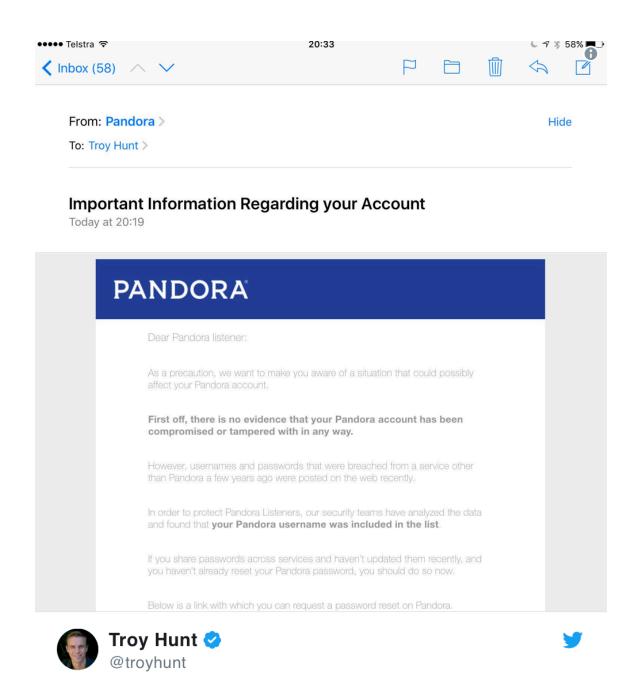
To make sure you continue having the best experience possible on LinkedIn, we're regularly monitoring our site and the Internet to keep your account information safe.

We've recently noticed a potential risk to your LinkedIn account coming from outside LinkedIn and just to be safe, we've locked your account for now. You'll need to reset your password in order to unlock your account. Here's how:

- 1. Go to the LinkedIn website.
- 2. Next to the password field, click the "Forgot your password" link, and enter your email address.
- 3. You'll get an email from LinkedIn asking you to click a link that will help you reset your password.
- 4. Once you've reset your password, a confirmation email will be sent to the confirmed email addresses on your account.

Thanks for helping us keep your account safe, The LinkedIn Team

The premise of companies accessing data breaches in order to protect customers has really taken off and frankly, sometimes I think it even goes too far. I received this one myself from Pandora who merely found my *email address* in another data breach:



Just got this from @pandora_radio, sign of the times:

♡ 15 12:34 AM - Jun 28, 2016 · Gold Coast, Queensland

The point is that there are ways to use this data *for good* and what we've seen over the last year and a bit is the bad players dropping off whilst the good players gained prominence. What that means for the industry is "a rising tide lifting all boats"; it's becoming more legitimate for all those doing the right thing with the data. Let me shift attention back to HIBP because there's been a heap of

other things happen over the last year that have really helped with the legitimisation the title of this post speaks about.

Breached Sites Have Been Embracing HIBP

One of the things that's really pleased me is the way breached sites have embraced HIBP after they've suffered a security incident. A great example of this is the self-submission of their breached data. On 3 separate occasions now, services that have suffered a data breach have reached out and said "we'd like our members to be able to confirm they've been impacted by searching HIBP". TruckersMP first did this in Feb 2 years ago, Ethereum followed in December 2016 as did biohack.me in August last year. They all recognised that HIBP is there to help victims of data breaches after things go wrong and willingly offered a copy of the data that was now in public circulation.

Then there was this one from Daily Motion in August:

We would like to stress that HIBP is a trusted website that enables its users to verify, in a responsible wather their account was part of any of the data leaks contained in its database. It does not disclose user passwords either in clear text or in their hashed form.





I'm very pleased to see @dailymotion reference @haveibeenpwned in this fashion after I loaded their data breach press.dailymotion.com/archives/1515

○ 27 2:46 PM - Aug 10, 2017

Their data <u>first turned up on LeakedSource the year before</u> (I suspect the original attacker was paid for it, hence it appearing there before anywhere else), so the data breach itself wasn't a surprise to them, but obviously once it

appeared on HIBP the incident received more exposure again.

Oftentimes, the first a company knows of a data breach is when I send them their data. In some cases, this really rattles the organisation, particularly those that are less well-equipped to deal with these incidents (i.e. not as tech savvy). But increasingly, I'm finding the engagement with hacked companies is being well-received, for example after the Disqus disclosure in October:

Thank you to Troy Hunt for initially alerting us of this.

And the We Heart It breach just a few weeks later (I referred them to the Disqus disclosure measure as an example of best practice hence the similarity in their messaging):

Thank you to Troy Hunt for initially alerting us of this.

Or it may even just be a little reference and a link per <u>Kickstarter's update</u> breach notice:

Update, October 6, 2017: Some of our customers are hearing about our 2014 security breach today from a breach notification service. A quick recap: Once we learned of this problem in 2014, we closed the breach, emailed all of our customers, and posted an alert encouraging everyone to reset their passwords. We've invalidated any passwords that weren't changed at the time. Since 2014 we've strengthened our security measures, adding features like two-factor authentication and the ability to see where your account has been accessed. Both of those are available through Account Settings.

There's no new information today that changes what we shared in 2014. Our original post is below.

Engagement with these organisations may not necessarily always result in them giving a hat-tip to myself or HIBP, but the experiences I've had with many of them have not only led to public disclosure, but also resulted in some very good

communication of the incident. For example, the imgur breach in November and the Ancestry data breach in December. The point is that the very organisations I'd once feared would react badly to a presence on HIBP are responding in quite the opposite way. Of course, nobody ever *wants* to have their logo on the who's been pwned page, but I'm finding organisations increasingly accepting of the fact that data breaches happen and they're simply getting on with the job of managing the aftermath in a responsible fashion.

But it's not just organisations that have already been pwned that are giving HIBP a shout-out, let me share some more proactive examples.

HIBP is Becoming the "Go-To" Resource for Protecting Accounts

Last month, I noticed this piece pop up on EVE Online:

A good way to see if your email has possibly been compromised in a data breach is to go to:

https://haveibeenpwned.com/

...and check there. There is also a lot of interesting information about various data breaches they have collected on the page. Many of the breaches listed have had a direct impact on our situation. For example, anyone using a yahoo email for EVE Online would be well advised to secure the email to avoid unwanted attention from thieves looking for stuff to steal.

I love this because it's proactive; it's encouraging people to make a behavioural change by seeing first-hand how extensively they've already been pwned. I saw the same thing again from Epic Games just a couple of weeks ago with the release of their Fortnite blockbuster:

HOW DO I KNOW IF I'M AT RISK?

0

There is a fantastic web service Have I Been Pwned that will let you search your email address and determine if it has been part of any data breaches. If it has, you should assume that the password associated with that service is public knowledge and change all accounts that use it (not just your Epic account!).

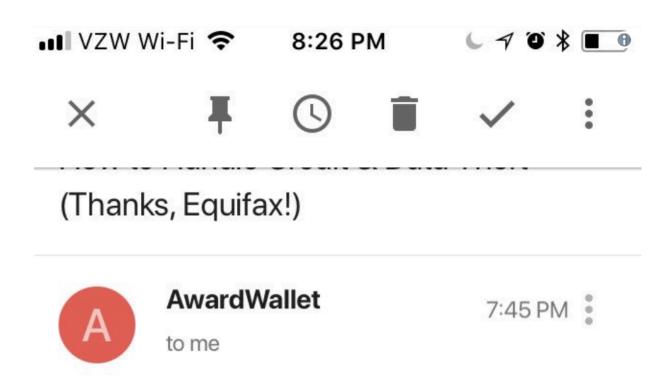
Even if your account information hasn't been publicly identified as leaked, it's possible that it may be leaked in the future, so there are steps that you can do to help protect yourself against that. You can start by signing up for the Have I Been Pwned notification service so you're immediately alerted if your email is ever included in future dumps.

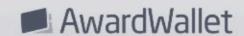




Sage advice by @FortniteGame! epicgames.com/fortnite/en-US...

Then there was this nice plug from AwardWallet back in November:





Data and credit theft have become too much of a norm solutions offered by the impacted businesses just don't

Advertiser Disclosure: The credit card offers that appear on AwardWallet are from credit card companies from which compensation. This compensation may impact how and where products appear on AwardWallet (including, for example they appear). AwardWallet does not include all credit card companies or all available credit card offers





- @troyhunt just got this from @AwardWallet for @haveibeenpwned
- O 4:27 PM Nov 28, 2017

Back in September, I saw the same again from Deliveroo

There are excellent free tools which our customers may use to help discover if they are at risk — for instance Troy Hunt's Have I Been Pwned? website; and we recommend use of such tools as an aid to password security.

And one that's come as a real surprise - I've heard many similar examples of the following advice from Netflix where an operator recommends HIBP during a support call:





@troyhunt Dealing with some Netflix account issues and the support rep directed me to haveibeenpwned.com. Maybe you can hit them up for a free account or your own movie. :-)

2:07 PM - Feb 5, 2018



Incidentally, it's reasons like the Netflix example which demonstrate the value of keeping this data publicly searchable, namely that it helps support staff establish possible sources of account takeover. I touched on this in <u>my September piece</u> on the ethics of running a data breach search service.

Recommendations for checking HIBP can come from places I never expected, for example <u>German company Stiftung Warentest</u>:



Incredible 5 billion stolen records from known services such as dropbox include the databases of two pages where you can check if your email and password combination has fallen victim to a hack.

The site http://haveibeenpwned.com has created a Microsoft employee. It includes many hacks from international us services. For example, if you have an account with Dropbox, Linkedin or badoo for longer, you will probably find his access in Simply Enter E-mail address and the results will appear immediately.

Whilst I may not have previously heard of them, apparently their opinion carries some weight:





Replying to @the_jannis @troyhunt

They're considered to have absolute integrity and is trusted by everyone here but conspiracy theorists, including the government afaik.

♡ 1 6:22 PM - Oct 1, 2017

Or from very familiar names, such as Google:

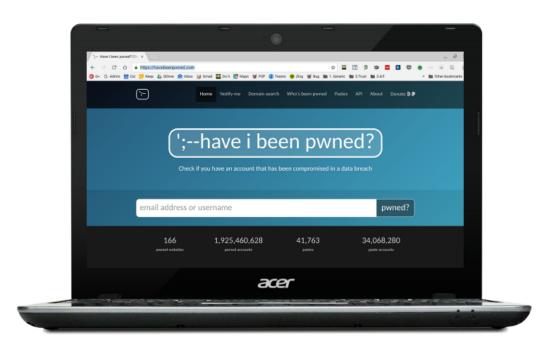




0

Nice to see @GoogleForEdu recommending @haveibeenpwned by @troyhunt to schools on @edugeek - awscdn.cdngeek.com/fls/spon/googl...

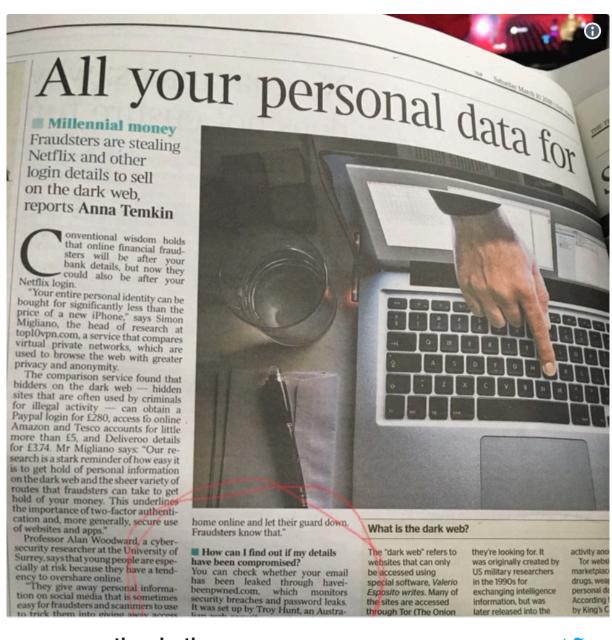
○ 6 3:20 AM - Dec 11, 2017



Google for Education

Consider why organisations like the ones above do this: they have to deal with account takeovers every single day - it's a massive issue. It's in their best interests to drive more positive security hygiene amongst their members and evidently, having customers check HIBP for breach exposure helps them do this.

Sometimes, endorsement even extends through to the real media!





mathewbutler

@mathewbutler

@troyhunt just to let you know that you and HIBP get a positive mention in the UK press

○ 8 1:16 AM - Mar 10, 2018

And a big "thank you" to that organisation for causing me to register yet another variant of the HIBP domain name to ensure that havei-beenpwned.com is fully

functional!

Apparently, HIBP is even getting mentions at Harvard these days:





Was excited when 'have I been pwned?' & @troyhunt was referenced in class at Harvard. Told the prof how I used to spend summers working in his office at Pfizer making binders for my dad!

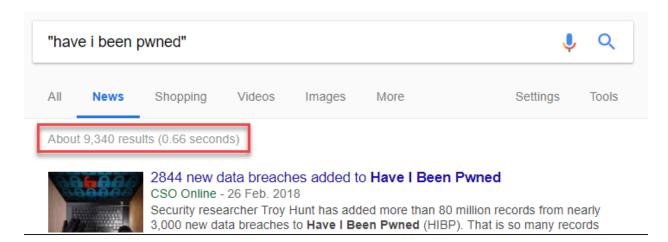


(Fun side story: Arjun's dad was my boss at Pfizer for about 14 years, must have been a weird coincidence when he heard HIBP mentioned!)

But as much as HIBP has received some great plugs by companies recommending people use it, it's the media that's generated the most attention.

Press

I actually used to maintain <u>a page listing major media pieces</u>, but the whole thing got too unwieldy as the press mounted. That's a link to an archive.org copy of it (I've since removed the page), I tend to just link people to <u>a Google news search result</u> these days:



I got to thinking about the press again this week after HIBP popped up on a Belgian TV show:

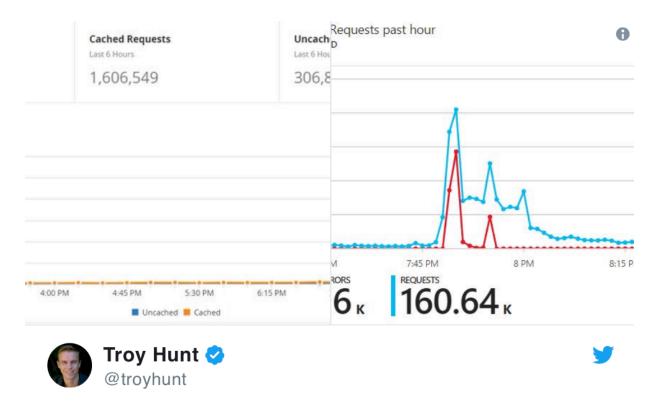






Tonight @haveibeenpwned was featured on Belgian TV @opVIER, but I wonder if @troyhunt will notice a spike in traffic $\ensuremath{\mathfrak{e}}$

There's a heap of similar examples to this, perhaps the one which made me most think about how I deal with the sudden influx of traffic was The Martin Lewis Money Show in the UK which ultimately led to this:

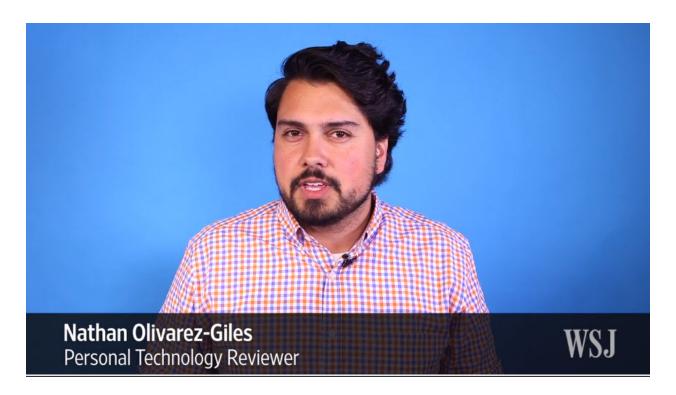


So @haveibeenpwned just copped a massive sudden spike of traffic sent faster than Azure could scale. The Money Show? troyhunt.com/brief-lessons-...

♡ 21 12:28 AM - May 15, 2017 · Gold Coast, Queensland

It's a fascinating scaling problem to deal with: how do you handle a 100x change in traffic volumes over a period of 60 seconds whilst also minimising infrastructure costs? I talk about it in that blog post and have since made some other big changes, especially to the aggressiveness with which Cloudflare caches content.

One of the great things the media has done for HIBP is to put it out in front of everyday people, that is folks who may not live and breathe tech like (probably) you and me. Of course, the media can totally misrepresent the mechanics of data breaches and how they actually occur (<u>as I've lamented before</u>), but occasionally, pieces like this one from the Wall Street Journal really nails it:



I especially like the focus on trustworthiness, plus of course the general good advice that an outlet like this can put out in front of normal folks. Mind you, those same media companies struggling with the name have caused me to register some rather odd domains including haveibeenprawned.com and haveibeenprawned.com, thank you very much.

But one of my favourites is one targeted more towards us tech people, and it's this one from WIRED:







Want to know if you've been hacked? @troyhunt has all the details wired.uk/NGmwcn

◯ 34 9:29 PM - Aug 8, 2016

I remember doing that photo shoot with them a couple of years ago, standing around in the rain in London whilst struggling with a cold and almost no voice. But hey, the pics came out great and I actually have a page from the *real* print WIRED mag framed on my wall now.

In gathering these references over the last 6 months or so, there was one particular source which popped up over and over again that really surprised me - the police.

Law Enforcement Has Been Extensively Recommending HIBP

Remember how I started this post by referring to all the illegal activity which led to HIBP even being necessary in the first place? You can imagine how tweets like this initially came as a bit of a surprise:



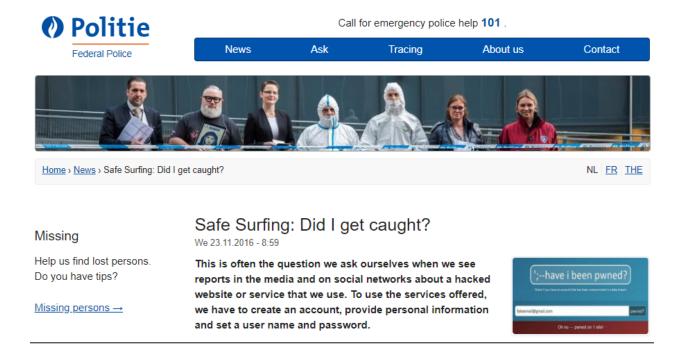




A shout out for @haveibeenpwned by @troyhunt at the Cyber Security Summit today. Held at @LancsPolice by @LanpacLtd and presentation by @TITANROCU!

♥ 12 2:46 AM - Nov 23, 2017

In fact, police forces all over the world have been publicly promoting HIBP, for example the Belgian federal police (Google translated for non-Dutch speakers):



And whilst I'm translating things from Dutch, <u>here's another one from the Netherlands police</u>:

Privacy

If the police roll up such a network, that still does not have all the cold to be out of the sky, warns De Milde: "There is always the possibility that your data have now been sold to others. And remember: even if you do not have data in the police file, there may still be something going on. Perhaps your data has been stolen by a botnet that we have not yet detected. Regularly replacing passwords remains a necessity and also look at check sites such as haveibeenpwned.com from time to time . For reasons of privacy, the police may not disclose confiscated data. That is why we are unfortunately not allowed to actively inform affected citizens themselves. Offering them the opportunity to consult this new database will fortunately enable us to be of service to our citizens and to reduce such crime with them.'

(Ok, we disagree on the regular rotation of passwords, but it's a nice shout-out all the same.)

But it's back in the UK again where law enforcement has been a regular supporter of HIBP via a number of shout-outs over recent months. For example, via the Police Service of Northern Ireland:







Looking to see if your email address has ever been compromised? Why not pay a visit to haveibeenpwned.com.

Remember always use a strong separate password for your email account. #PSNICyberProtect @CyberProtectUK @cyberawaregov @actionfrauduk

○ 37 5:30 AM - Mar 13, 2018

Back in England, the Leicester Cyber Aware account (and their dogs) recognise HIBP's role in keeping people safe:





#FF These guys ...to keep you safe @waynedenner @GlosSaferCyber @TakeFive @UK_SIC @kentpolicecyber @HP_Cyber @MFF_Forum @haveibeenpwned @EC3Europol @TheParentsZone @CifasUK @AgainstScams @GetSafeOnline @CyberProtectUK @NottsFraudCops @thebreckfound @gcluley @WMP_ECU \$\times\$ 13 5:23 AM - Dec 8, 2017

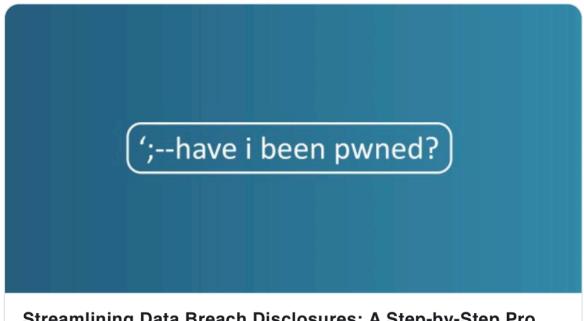
A bit further south and the Devon and Cornwall Police's Cyber Protect team feels the same way:



"Many well-intentioned people simply give up and don't report serious security incidents when the effort is too high or the risk is too great. That has to change"

Well said! @troyhunt #databreach #hacking #hibp #CyberSecurity troyhunt.com/streamlining-d...

○ 17 11:35 PM - Jan 14, 2018



Streamlining Data Breach Disclosures: A Step-by-Step Pro...

I don't know how many data breaches I'm sitting on that I'm yet to process. 100? 200? It's hard to tell because often I'm sent troyhunt.com

Over to Kent and it's the Police Cyber Crime Unit's turn:

A



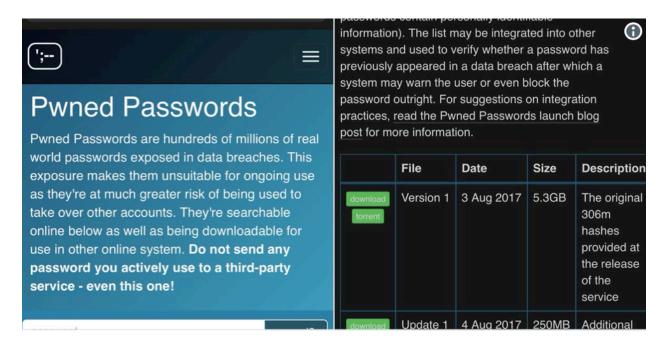




#CyberSecurityawarenessmonth, Check to see if your email address has been compromised? haveibeenpwned.com

♥ 5 3:30 AM - Oct 4, 2017

Even Police Officer Tony Murray recently gave Pwned Passwords a plug and offered some very good advice whilst doing so:





Tony Murray - National Protect Officer, NCO @CityPoliceTell2



♠ONLY check active passwords via the #DOWNLOADED list!

You have strong passwords, you use different #passwords for different accounts AND YOU could still be compromised.

Are your passwords already part of the 306 million already known?haveibeenpwned.com/Passwords #Tell2

○ 47 3:51 AM - Jan 29, 2018

All of these came as a surprise and getting back to the original context of this post - "the legitimisation of HIBP" - you can see why I value them. Having law enforcement speak in glowing terms has been *enormously* encouraging. Let me now take that one step further and talk about government.

HIBP is Increasingly Recognised by Governments Around the World

I'm going to talk about some government exposure that will be pretty familiar to many people shortly, but it's the unexpected stuff that, well, I just never saw coming. For example, the Estonian CERT advising people to check HIBP:

CERT advises all people who want to check whether their passwords have leaked to visit the website haveibeenpwned.com. CERT also continues to advise people to regularly change their account passwords, because when the password of one website leaks and a person uses the same password on other sites, criminals will get access to other accounts as well.

Much closer to home for me, our local Aussie Government recently gave HIBP a shout-out <u>via their Stay Smart Online initiative</u>:

Find out if your email address has been breached

To find out if your email address has been published in a data breach, go to $\frac{\text{HavelBeenPwned}}{\text{Model}}$ and follow the prompts.

But there were 2 especially important recent events tied to government and I want to spend a bit of time explaining the significance of both. The first one was this:



This was my testimony to US Congress in November (there's a video of it in that link). Some people actually joked in advance of this that the invitation was a means of getting me over to the US so that I could subsequently be locked up for sitting on billions of records of breached data! In reality, quite the opposite happened: I sat in front of law-makers and talked about this industry I've found myself in, including the relevance of HIBP. I never saw that coming and it will stick with me for life as both a momentous occasion in my own career and indeed a milestone in the history of HIBP.

More recently, I was especially proud to see something I'd been working on for the past 9 months finally come to fruition:





We're excited for the opportunity to work alongside Troy and we're looking forward to trailing the use of his service in the coming weeks to help alert UK Government departments if their users have potentially compromised credentials

Troy Hunt 🤣 @troyhunt

I'm proud to announce that @haveibeenpwned is now being being used to monitor UK and Australian government domains nationwide. I've been working with the @NCSC and ACSC to ensure they have a model where they can get fast access to their own information troyhunt.com/the-uk-and-aus...

○ 187 10:18 PM - Mar 1, 2018



As I wrote earlier this month, both the NCSC in the UK and the ACSC in Australia are now using HIBP to monitor their government domains. This was enormously important to me on many levels; it was obviously recognition from the respective governments that HIBP has a role to play in protecting their people, but it was especially poignant to me that both governments were also happy to acknowledge it publicly. That's a really big deal in terms of the whole legitimisation piece and certainly it was something I was especially conscious of as the arrangement fell into place. These two governments won't be the last either - I'm presently in discussion with multiple other departments from different parts of the world and I hope to be able to share the outcome of that shortly too.

Summary

When I started HIBP back in late 2013, I never envisaged any of what you've read above. Over the last 4 and a bit years, there's certainly been some ups and downs in terms of how comfortable I've felt with the legitimacy of the service and obviously I'm now exceptionally happy with where it sits today. It's where it is due to a combination of good luck and good management; I've been fortunate with the timing in the industry in terms of the prevalence of data breaches, but I've also been exceptionally cautious with how I've positioned HIBP, how I've engaged with corporations and governments and indeed the moral compass I've run it by.

I'll have more to share on the HIBP roadmap in the near future, this post was really just an opportunity for me to take a moment and reflect on where things stand today. A big thanks to everyone who has supported both the project and myself to help get it to this point!

Comments

Given the attitude of "shoot the messenger" that has historically been a major feature of the security industry, I am of the opinion that your decision to set up HIBP in the first place was incredibly brave. You've acted ethically and above-board throughout, but that hasn't always saved others, and there's always the danger with this kind of thing of being put in a compromising position that can be difficult to get out of.

It looks to me -- certainly over the last few months -- that you've reached a tipping point where the service is sufficiently legitimised that you can relax a bit about all of that. This post bears that out.

I'm really glad the HIBP exists. Well done.

Troy: I've certainly been conscious of that and part of my "insurance" against someone wanting to shoot the messenger is the transparency with which I've

run the service. I'm crystal clear about all the decisions I make and the engagements I have with organisations so if anyone ever wants to play that game, then it too will be played out very publicly and they'll be judged accordingly.

Having said that, I think that much of the reason I've never received a legal threat or had to deal with angry people is that I've always been open and receptive. There's certainly been organisations that have contacted me and had I behaved differently, they may have gotten quite upset. It's amazing how far simply being nice to people gets you in these situations:)

I have to ask a question that's been nagging me for a while... Do you ever regret the name of the site, given it's likely most folk don't know how to pronounce it and might think it's 'weird'?:D

Troy: No, never, in fact it's proven enormously good at generating discussion. People want to know how to pronounce it, what it means, where I got it from etc. Never once have I seen evidence that the strange name has been detrimental, at least not beyond causing me to register a bunch of mis-typed domains!

I demo'd HIBP to a bunch of my users in a security awareness training session the other day. The slide title I used was My Little Pwny

I've thought about this before, and this post made me wonder it again. How big is the "bus factor" in HIBP? I love what you're doing and how you're doing it, but what would happen to it if you where to get "hit by a bus" figuratively speaking (aka unable to work on it for whatever reason)?

Troy: I'm very conscious of that and I do have some things in the works around it. But frankly, if I was to get hit by a bus, it's not exactly like there's immediate and significant consequences for those using the service. But yeah, I'm working on it!

I think your post highlights that there's a clear need for a service such as HIBP. However, in terms of the future I suppose my main concerns as a user are as follows:

- Should a single individual be responsible for running such a service?
- Should such a service be "for profit" or not?

In terms of point 1, I absolutely trust Troy Hunt as an individual. I think you've done a fantastic job creating a useful, necessary service in such a way that it's trustworthy, transparent and puts users privacy first. But what happens if you get hit by a bus tomorrow? What happens to the service? I'm sure others will take up the mantle, but who has that trustworthyness that you do? Do you have a plan in place for a successor in such an event? Is the answer to spin off HIBP into a separate entity that's initially ran/owned by yourself, but has the workings of a team that can take over in such an event?

In terms of point 2, you've put a significant amount of your own time and money into the project. I can only applaud you for such a thing and I'm aware that you're almost certainly not going to get rich from it. You're careful with how you write your API's and utilise resources so that costs are minimised and you get some support from other vendors, like cloudflare - but ultimately, you've got to eat, right? I don't mean to presume about your own financial situation but tying into point one, would any kind of successor be tempted to run HIBP for-profit?

It seems to me that perhaps the future would involve turning the project into some kind of not-for-profit entity?

Troy: The reality is that many individuals and corporations alike run similar

services and frankly, I'm more interested in the ethics that guide them than I am about someone having access to the data. And see my reply to <u>Olle Kelderman</u> regarding the bus situation.

To date, HIBP has been something I do in my spare time, perhaps consuming about 20% of my overall week. It's helped enormously with the courses I create, the blog posts I write and the talks I deliver because it gives me enormously valuable insights into the world of data breaches. Those activities are very much commercial (in one way or another) and that keeps me well-fed. I don't see any reason to change the present structure at this time but I do have thoughts on where it might go in the future that will ensure sustainability as well as unlock a bunch of untapped potential. For now though, everything is in good balance:)

Troy, you probably know you can count on support from your users when and if you decide to establish some kind of org to handle HIBP's maintenance, finances, and succession plan.

When and if you're ready for that, just ask.

Epilogue

I'm writing this epilogue on the 27th of April 2021, mere hours after publishing a blog post about how the FBI just gave me 4.3 million email addresses from the Emotet malware. When I read back through the legitimisation blog post of 3 years ago, I can't help but feel pretty damn proud of how far this thing has come. It's no longer about trying to convince the public I'm running a legit service because even law enforcement agencies recommend it, I've literally just loaded in millions of records of PII from the FB-freakin-I! That's cool. Really cool



What I feel most proud about isn't that this is somehow a reflection of my glorious intelligence (I intentionally made that sound stupid because I find it a

frankly stupid concept), rather that something I created and so carefully nurtured could gain the trust of the most well-known law enforcement agency in the world. Throughout our discussions in the lead-up to this blog post and publishing the millions of records into HIBP, I kept waiting for the point at which they'd think twice about providing all this personal data to me. Sure, it's "only" email addresses, but in an increasingly privacy-aware era, I wouldn't have been surprised if a lawyer somewhere had squashed the idea. But what would that have achieved? The individual would still be compromised, and the only difference of any consequence is that they wouldn't know it.

And this, my friends, is what I see as the next frontier: breaches happen, personal data gets exposed, so now what? Do we tell everyone they should delete any copies they have of the data, or do we seek out the most impactful thing we can possibly do now that the data is out there? (Fun fact: in September 2021 following the breach of Epik, a registrar and host popular with extreme right-wing services, their CEO Rob Monster recorded a video saying the breached data was cursed and anyone who had it should delete it. There is no evidence this strategy worked.) I've had precisely this discussion with higher-ups in GDPR and regulators are still at this impasse of somehow reconciling the intention for data to only be used in the way in which its owner provides it whilst dealing with the reality that it's being used in all sorts of other ways counter to what we'd like. I hope a future blog post on the legitimisation of Have I Been Pwned can answer the "now what?" question and this original post and collaboration with the FBI are just 2 steps along that journey.

PWND PASSWORDS IN PRACTICE: REAL WORLD EXAMPLES OF BLOCKING THE WORST PASSWORDS

If you've ever built a little personal project yourself and put it out there for the world to use, you'll be familiar with the feeling of wondering if anyone is actually using the thing. Like me, you might have poured your heart and soul into it, sitting in front of the keyboard at all hours and you've done it not because it's going to make you a gazillionaire, but for the love of it. If nobody was ever going to use the thing, you'd probably still have done it anyway because it scratched an itch somewhere. A desire to create something or even just solve a problem for yourself. However...

Seeing your baby make a bigger impact on the world is amazing. That's how I felt when writing this post: "hey, look, I've built something useful that's actually helping people!" I wanted to highlight that in this post because I wanted more people to use it. I wanted to show that there were some big names and big use cases and hey, maybe you could use it too. This post felt like it was the coming of age for Pwned Passwords.

29 MAY 2018

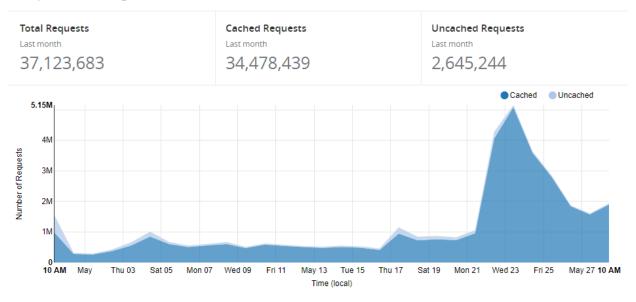
Back in August, I pushed out a service as part of <u>Have I Been Pwned</u> (HIBP) to help organisations block bad passwords from their online things. I called it "Pwned Passwords" and <u>released 320M of them from real-world data breaches</u> via both a downloadable file and an online service. This was in response to <u>NIST's Digital Identity Guidelines</u> and in particular, the

following recommendation:

When processing requests to establish and change memorized secrets, verifiers SHALL compare the prospective secrets against a list that contains values known to be commonly-used, expected, or compromised. For example, the list MAY include, but is not limited to: Passwords obtained from previous breach corpuses.

Seen a password in a data breach before? Then now it's a *pwned* password and per NIST, you really don't want to be letting your customers use it any more. I followed up the first version with <u>version 2</u>, <u>complete with just over half a billion passwords</u>. But the *really* cool bit was the k-anonymity model devised by Cloudflare which I talk about in that blog post. That really started getting the service traction, but it wasn't until last week that things *really* started to fire:

Requests Through Cloudflare



That up-tick on the 17th and then the *really* sizeable one on the 22nd are due to a few big players making really good use of the service. I want to detail those use-cases here because I'm always getting asked by people how the service is being used. So here it is - including some inside stories - Pwned Passwords in practice!

1Password

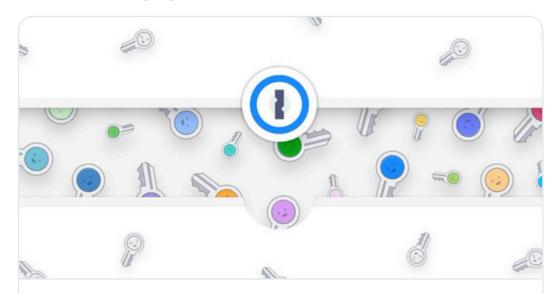
7 years ago now, I realised that the only secure password is the one you can't remember and from that day forward, I've been using 1Password exclusively as my password manager. When I released version 2 of Pwned Passwords, out of the blue they built it into their product. This wasn't some big effort on their behalf which took lots of planning either, they literally did it overnight:



Troy Hunt 🤣 @troyhunt · Feb 22, 2018



Hey, you know what would be cool? If @1Password was to integrate with my newly released Pwned Passwords k-Anonymity model so you could securely check your exposure against the service (it'd have to be opt in, of course). Oh wow-look at this! blog.agilebits.com/2018/02/22/fin...



Finding Pwned Passwords with 1Password - AgileBit...

1Password now includes a proof of concept to allow you to check if your passwords have been leaked on the internet blog.agilebits.com

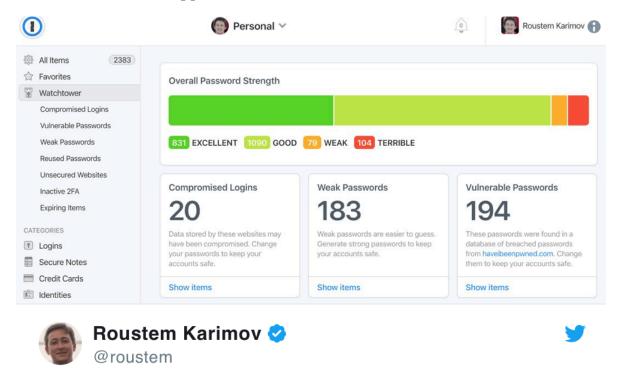


Troy Hunt <a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O<a>O

I'm *so* impressed with what they've done here; I launched this service only 27 hours ago and they've already pushed this out. They had no prior knowledge I was doing this, they just got hands on tools right away and made it happen. That's awesome.

○ 507 12:15 PM - Feb 22, 2018

That was impressive and they gave people the ability to check any individual password against the online Pwned Passwords service (it was also part of the reason I ended up partnering with them on HIBP). It used the k-anonymity model to ensure the original password wasn't redistributed, but it was a one-by-one effort. Until this happened:



If you use 1Password account you now have a brand new Watchtower integrated with @haveibeenpwned API. Thank you, @troyhunt

Also, looks like I have to update some passwords (\$\sigma\$) 227 7:19 AM - May 3, 2018

They built in the ability to check *your entire set of passwords* against Pwned Passwords in a single action. That was via the web-based version of the tool and they followed that up last week with the launch of 1Password 7 for Mac, including Pwned Passwords as a first class citizen of the desktop app:

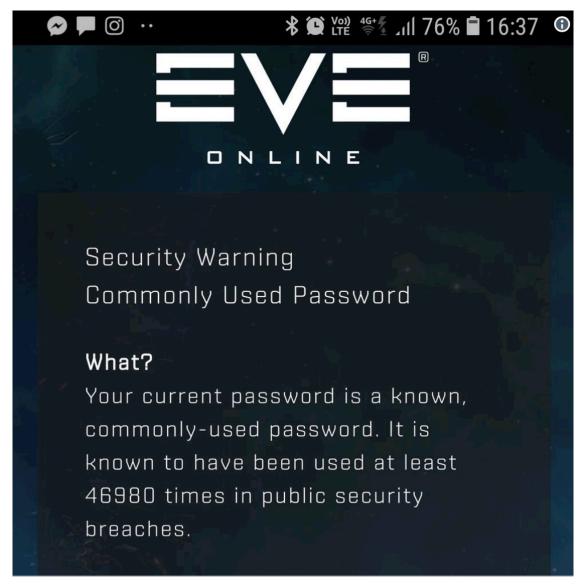
Watchtower integrates with Troy Hunt's haveibeenpwned.com service to see if any of your logins are vulnerable. 1Password securely checks your items against a collection of breached passwords (over 500 million and counting) and notifies you to change them.



This week will see the launch of 1Password 7 for Windows so I'll be jumping on that one pretty promptly. I've been told the Pwned Passwords integration will come shortly after launch, but it's available to everyone already via the web version if you just can't wait.

EVE Online

Last month, I got the first indication that <u>EVE Online</u> - the massive online multiplayer RPG - was planning to implement Pwned Passwords:

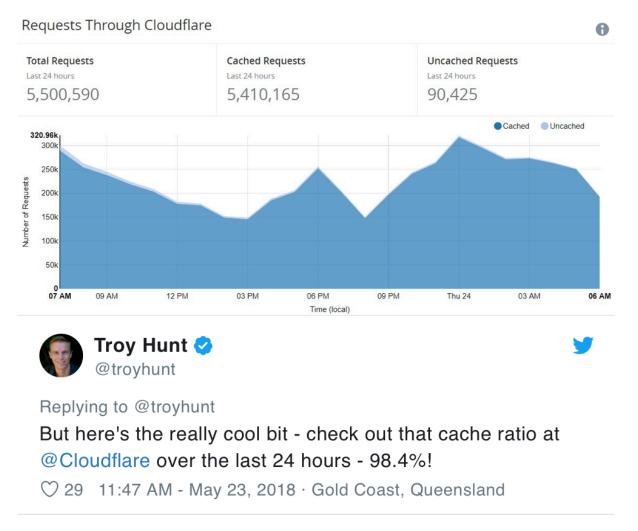




WIP: Helping our @EveOnline players to be aware if their passwords are on a list of known compromised passwords. Thanks @haveibeenpwned! CC: @troyhunt #tweetfleet #security #workinprogress

Stefán and I ended up talking quite a bit, especially around optimisations to the Cloudflare caching implementation to ensure it was going to be *super*-fast for

them. Working with <u>Junade Ali from Cloudflare</u> who devised the original k-anonymity model, we got the cache-hit ratio *way* up:



I've actually seen it round to 99% before as well but hey, now I'm splitting hairs! The point is that almost every single request to the service is now hitting one of Cloudflare's 150+ edge nodes around the world and returning the result in what is usually no more than low double-digit milliseconds.

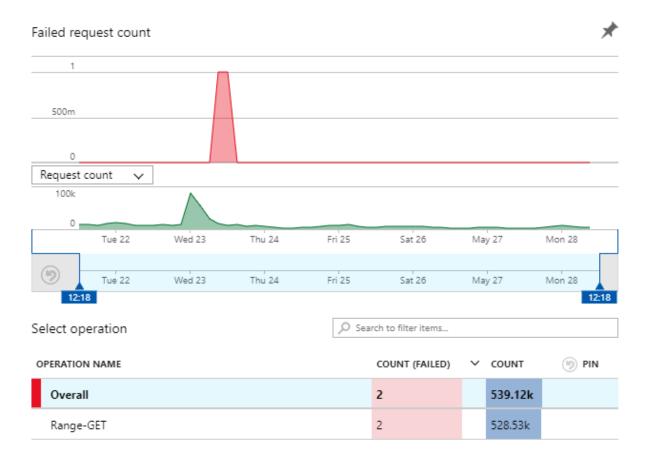
Because I really wanted to share some real-world info on how the service is being used, I asked Stefán if he'd mind me publicising some stats which he kindly agreed to. Keep in mind that at the time of writing, the Pwned Passwords check is only on the login, it'll shortly be rolled out to the registration and change password features as well:

- 1.EVE online is making 40k requests per day to the API
- 2. The median response time for the service is 18ms
- 3. The 99th percentile response time is 930ms
- 4.On launch of the service, 18.6% of passwords were found to be previously pwned
- 5. Several weeks later, that's now dropped to 17.3%
- 6.They've also seen an up-tick in people updating their passwords and enabling 2FA

Stefán also shared some info on failure rates and I'm just going to quote him directly here:

Errors almost nonexistent, usually caused by network blips on our side. Out of almost 300k requests we've seen only 100 errors in the last week. That's a 99.96% successful request ratio and we can't for sure blame those 0.04% on the API as they could be on our side :)

I don't think that's entirely fair; I have actually seen 2 failed requests in the last week:



But depending on how you look at it, that's either a 99.999% success rate on Azure Function executions or a 99.99999% success rate on all calls made to the API because so many are returned by Cloudflare. Seven nines - I'm happy with that

Kogan

<u>Kogan</u> is one of our largest online retailers down here in Australia. Like most stores on the web, people log on, store personal account info and, of course, buy products. As of last week, Kogan is using Pwned Passwords to help protect those accounts:





Replying to @troyhunt @ubernostrum

Well it took a bit longer than we thought, but our signup and change password flows are now using the HIBP password API. kogan.com/au/password-sa... is linked to in the validation errors. We'll iterate over time. Happy to hear feedback too!

○ 4 7:51 PM - May 23, 2018





Password Safety - Kogan.com

Kogan.com is Australia's largest online retailer, with incredible deals on TVs, Phones, Tablets, Computers, Kitchen Appliances, Homewares,

kogan.com

By virtue of me already having an account on the site, testing this was an easy one:

Confirm Password

•••••

Please use a different password, the one you provided is not secure. To learn about password safety, visit kogan.com/au/password-safety

UPDATE

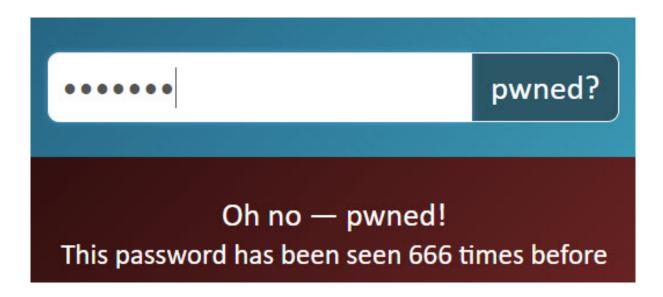
Pretty simple stuff, as is <u>the password safety page</u> they refer to. A site like Kogan's is used by the masses and we know empirically that your average person doesn't make the best choices when it comes to choosing passwords. How bad are some of those choices? This bad:



Today it's World Password Day: choose a word that's already in your heart. Like "Nutella", for example! #WorldPasswordDay #Nutella

○ 273 3:04 AM - May 3, 2018

Yes, that's a real account and yes, <u>people do actually use "Nutella" as their password:</u>



Except people on Kogan - they don't use Nutella anymore 😜

Okta

I actually had <u>Randall Degges from Okta</u> reach out a couple of weeks ago and mention he was building a little tool called PassProtect to check passwords against the API using a browser extension. That's cool, I get a lot of people emailing me about similar things, so I replied and moved on. And then Okta launched it and somehow snagged *a heap* of news headlines:



REVIEWS

NEWS

VIDEO HOW

SMART HOME

CARS

DEAL

SECURITY

Okta's Chrome plug-in tells you when hackers have your password

And yes, it does have a system to avoid leaking your password itself.

FORTUNE

TECH . TECHNOLOGY

New Chrome Extension Warns You of Stolen Passwords



Okta's PassProtect checks your passwords with 'Have I Been Pwned'





Home | U.K. | U.S. | News | World News | Sport | TV&Showbiz | Femail | Health | Science

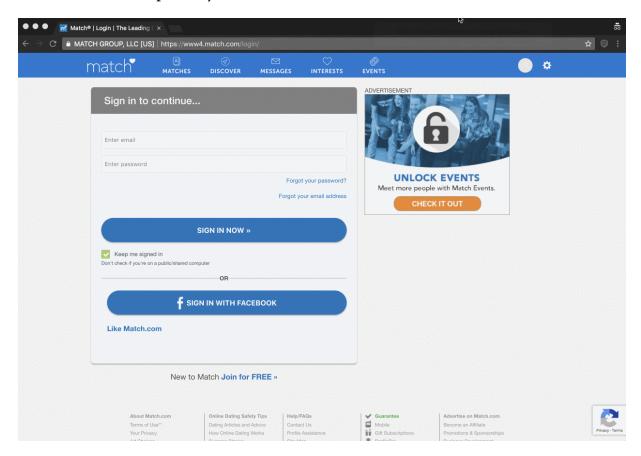
Latest Headlines | Facebook | YouTube | Google | eBay

Have hackers stolen YOUR password? New Google Chrome plugin will tell you if your details have been breached

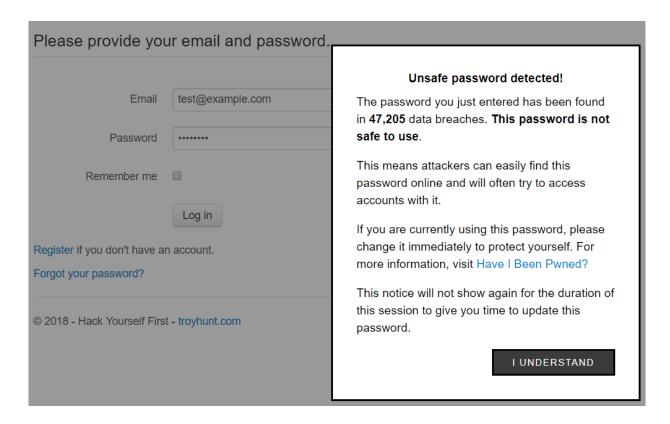
- A new Google Chrome extension tells users if their password has been leaked
- Cloud software firm Okta developed the Chrome extension, called PassProtect
- Once a user enters their password, the extension runs an encrypted version against the 'Have I Been Pwned' database to see if hackers may have it
- · If its been leaked, the extension will inform users that the password isn't safe

I love that this isn't just tech headlines either, it's consumer press like Fortune and the Daily Mail (do read the comments on that one, just for fun...) The point is that it gets the concept of how poor passwords are exploited out in front of the masses, and that has the potential to lead to very positive changes in our overall security posture as an online community.

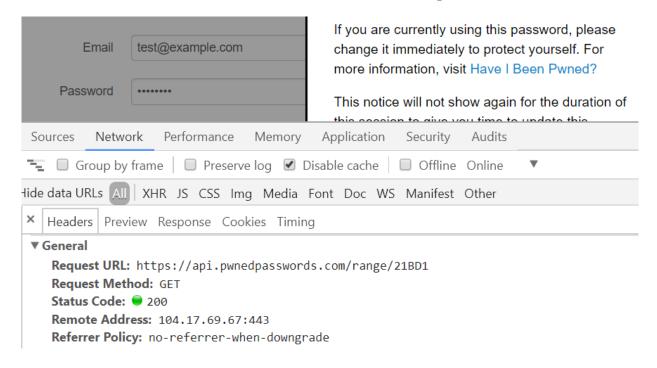
Okta wrote about their extension last week and they have a neat little demo GIF here that sums it up nicely too:



They've also got a <u>dedicated PassProtect website</u> and as you'll see there, they've made the extension both free and open source. I thought I'd give it a run on <u>the login page of my Hack Yourself First site</u> (a deliberately vulnerable site I use for training) and it worked beautifully:



And just in case you're curious, you can observe the call to the Pwned Passwords API in the browser dev tools once focus comes off the password field:



This is a really neat implementation by Okta with the extension simply looking

for changes in password fields. I don't mean "simply" in any derogatory form either, there's beauty in the simplicity and that's why they're (quite deservedly) getting such good press.

Other users

I wanted to highlight some of the biggest use cases via the organisations listed above, but I also want to acknowledge some of the multitude of others I've seen pop up just in the last week.

For example, Bittylicious:







In February 2018 we integrated with the excellent HavelBeenPwned service by @troyhunt to protect users from insecure passwords ow.ly/xFkP30k6RFK

♥ 4 12:15 AM - May 22, 2018

There's also Red Shield down in my corner of the world (ok, so New Zealand is a

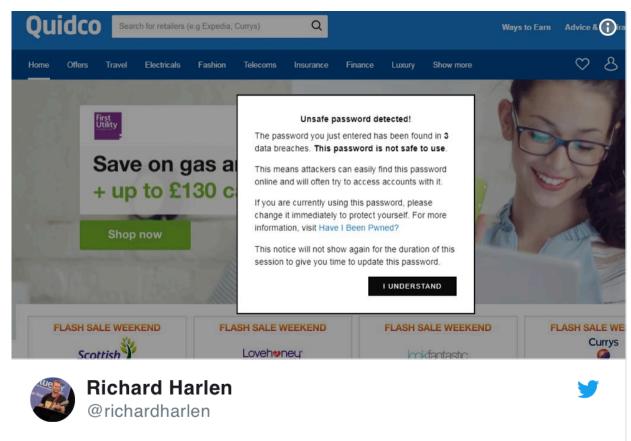
couple of thousand km away, but that's "close" down here!) who provide shielding services to websites. They recently began offering Pwned Passwords to their customers as part of their services to help protect websites:

ZERO	Home	About	API	Contact	
Register. Create a new account.					
You have entered a comm	non passwo	rd.			
Email	test23@	Dredshield.c	0		
Password					
Confirm password					

What I found interesting when Red Shield reached out recently was this comment:

It's much easier to communicate "Your password is a commonly used password" to users than complexity or entropy requirements

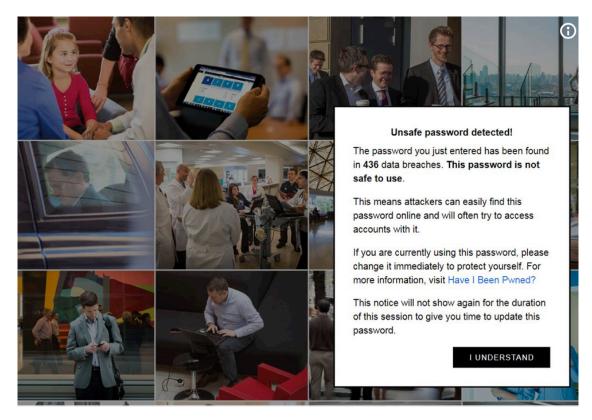
I wholeheartedly agree with this and as I've written before, <u>strength indicators</u> <u>help people make ill-informed choices</u>. Don't get me wrong - you still want a minimum bar to some degree (i.e. a min of 8 chars) - but we know that practices like character substitution or adding common punctuation to the end is extremely weak. Then there's the UK-based company <u>Quidco</u> who runs a cashback service:



Great to see Quidco.com has integrated with haveibeenpwned.com's password checking service! Time to change my password methinks. Thanks @troyhunt - I owe you a few beers by now I'm sure.

♥ 9 2:39 AM - May 25, 2018

Even the web interface for my Belgium mate John Opdenakker's mail account is using Pwned Passwords:







Just noticed that our webmail now has @troyhunt's pwned passwords integrated. This is really awesome

#Security

10 8:46 AM - May 28, 2018

And you know the really cool thing about those last 2 tweets? Neither organisation had to lift a finger because <u>that's actually Okta's PassProtect in action</u>. It's such a slick, integrated experience that both Richard and John didn't even realise the respective services hadn't done any work! That's cool.

And Finally...

Lastly, I just wanted to reiterate the message I provided in the launches of both V1 and V2 of Pwned Passwords: this is 100% free. Not free as in "if you're not

paying for a product, *you're* the product" either; there's no attribution requirement (I welcome it, but don't require it), you can do whatever you want with the downloadable data if you don't want to hit the API and if you *do* want to use the web service, there's no rate limit (quite the contrary as I've put a lot of effort into ensuring you can absolutely hammer it). There's absolutely no commercial angle from my side either; there's no "enterprise" version of Pwned Passwords, no up-sell and frankly, I've got no idea who's even using it beyond those who've explicitly told me. I do this with Cloudflare's support and because put simply, it's just a good thing for the web. I get a great deal of satisfaction out of building stuff that people love to use and given my access to passwords and having a platform to share this on, I've been able to make it successful and provide something genuinely useful. That is all .

Comments

Really nice to see this becoming more wide spread to non-tech circles as well.

However, the comments below the Daily Mail news leave me a bit sad and disillusioned at the same time. None of those folks seem to get the idea behind HIBP and PwnedPasswords. While I respect their skepticism about giving away passwords, this shows that it is difficult to sell the whole concept of PwnedPasswords to "ordinary" people. Most of them seem to have learned to keep their passwords secret (good!), so the idea of any "external" password checking will never seem to be a good idea to them.

Therefore integration into password managers and login services themselves seems to be the way to go! But this also means that these services better not say a thing about PwnedPasswords because it will only confuse people...

I really come to the conclusion that often tech ain't the problem. People are!

Troy: I understand the skepticism and you're right, integration with other

services is what really gets it traction. Fortunately, this means both those who understand it can use the likes of Okta's extension and those who don't still benefit when they go shopping at Kogan.

Hey Troy, on your last point, how are you financing the costs of running the servers? Or is this all being done on Cloudflare's free tier?

Troy: Cloudflare is providing their service which obviously absorbs the huge majority of requests. I'll need to check how it pans out this month, but last month the Azure Functions utilisation was *well* below the free threshold anyway ...

Is the free threshold reset every month?

Troy: Yep, but it would barely matter anyway given how cheap Functions are to run, especially fast ones. I'll share details once the next bill hits.

Jump on <u>azure.com</u>, and give their Azure Functions costing calculator a spin. You get a free quota of executions and utilization (a compound measure of CPU/RAM/per second, approximately). Developing something for a client recently and the entire API side is currently running free of charge, only costing is our database which is a measly \$5/mo or something.

Disclaimer: I'm not a Microsoft representative or anything, just a developer recently dropped in the deep end of Azure Cloud

Epilogue

I had to triple check the figures in this post to make sure I was understanding the time period properly; was that 37M million requests in an entire month? That's a single busy day now! In fact, at the time of writing, Pwned Passwords has passed 1 billion requests a month, which totally blows my mind. Yet I'm conscious that I'm sitting here now writing about that with the same amazement with which I wrote about 37M in a month back in 2018 and wondering how long it will be before I reflect on "only" 1 billion queries per month. That'll be cool

As you'll read later on, Pwned Passwords would later go through some far more significant changes that would make it both far more valuable and far more open and freely available than at the time of this post. I believe a huge part of what made that possible is it's simplicity; there are some super sophisticated credential stuffing defences out there that are very good at what they do, but with that comes complexity, cost and all sorts of other challenges (accessibility and maintaining user experience are but 2 examples).

So, here's my great insight: create a service that's simple, easy to use and free whilst providing value and you've got a winner!

WE'RE BAKING HAVE I BEEN PWNED INTO FIREFOX AND 1PASSWORD

I put a lot of thought into how to make HIBP the most effective tool it can be. Going back to that mantra I've repeated so often - "Helping people do good things after bad things have happened" - I was always interested in ways to amplify the reach. The service was tracking along well and had certainly exceeded every expectation I ever had of it, but what if it could do... more?

I had an existing relationship with Mozilla, not in a commercial sense but in a "hey, I'm in San Francisco, want me to come do a talk at Moz?" sense. Lots of discussions that started out very casually ultimately led to HIBP being integrated into Firefox. It was the same with 1Password with whom I'd already had the better part of 7 years plugging their product, not because it made me any money, but simply because I liked it ② So, the integration with these products came about completely organically, simply because it felt like the right thing to do at the time.

26 JUNE 2018

Pretty much every day, I get a reminder from someone about how little people know about their exposure in data breaches. Often, it's after someone has searched <u>Have I Been Pwned</u> (HIBP) and found themselves pwned somewhere or other. Frequently, it's some long-forgotten site they haven't even thought about in years and also frequently, the first people know of these incidents is via HIBP:

You've been pwned!

You signed up for notifications when your account was pwned in a data breach and unfortunately, it's happened. Here's what's known about the breach:

Email found:	@yale.edu
Breach:	Ticketfly
Date of breach:	31 May 2018
Number of accounts:	26,151,608
Compromised data:	Email addresses, Names, Phone numbers, Physical addresses
Description:	In May 2018, the website for the ticket distribution service





large @ticketfly data breach. thanks @troyhunt for the excellent @haveibeenpwned service that notifies users of #privacy disasters like this :) cbsnews.com/news/ticketfly...

○ 2 3:23 AM - Jun 4, 2018

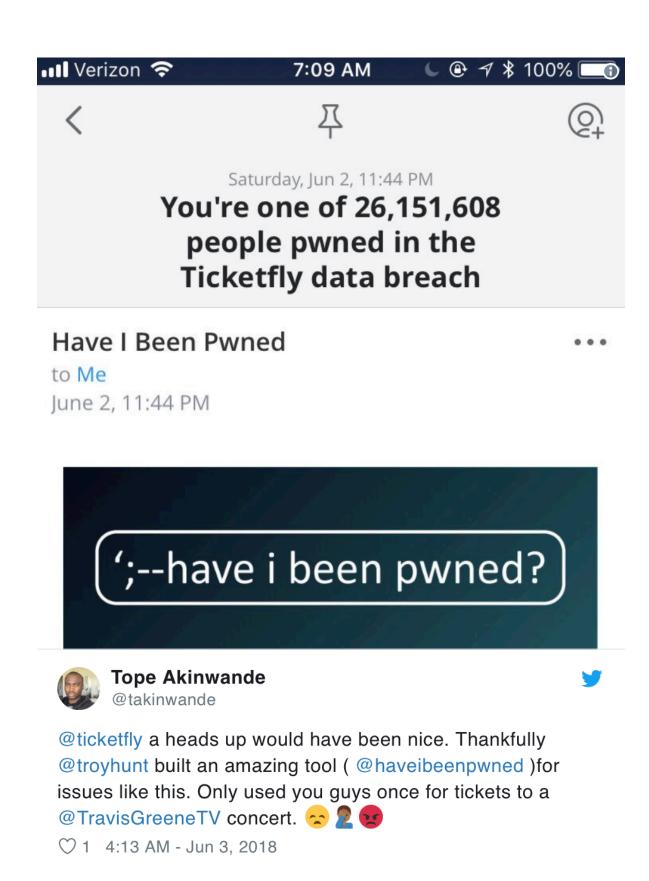




0

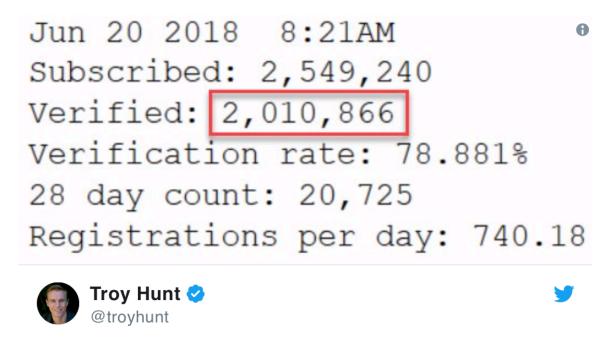
Well, that's annoying: @TicketFly data breach attacker publicly posted my info (along w 26MM others). I at least know about it, thx to @haveibeenpwned

○ 8:09 AM - Jun 3, 2018



In cases like Ticketfly, loading the data into HIBP meant notifying 105k of my

subscribers. That's out of a subscriber base that just recently ticked over the 2M million mark:



Wow, just realised @haveibeenpwned passed the 2 million *verified* subscribers mark whilst I've been travelling. That's amazing, never expected to see that!

○ 259 10:29 PM - Jun 19, 2018

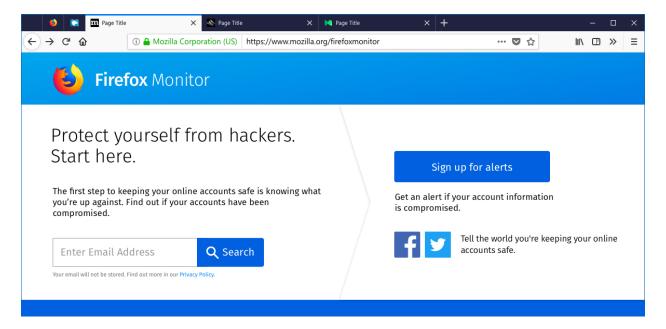
2 million is more than I ever expected, if I'm honest, but it's also only a tiny, tiny drop in the ocean. Of the 5.1 billion *records* that are in HIBP today, there's 3.1B unique email addresses. I'm reaching 0.06% of them via the notification service and not a whole lot more in terms of people coming to the site and doing an ad hoc search (usually 100k - 200k people a day). Don't get me wrong - I'm enormously happy and personally fulfilled by having been able to do even this - but clearly, I'm barely scratching the surface. However, that scope is about to expand dramatically via 2 new partnerships which I'm announcing today, starting with Firefox:

Mozilla and Firefox Monitor

Last November, there was much press about Mozilla integrating HIBP into Firefox. I was a bit surprised at the time as it was nothing more than their Breach Alerts feature which simply highlighted if the site being visited had previously been in a data breach (it draws this from the freely accessible breach API on HIBP). But the press picked up on some signals which indicated that in the long term, we had bigger plans than that and the whole thing got a heap of very positive attention. I ended up fielding a heap of media calls just on that one little feature - people loved the idea of HIBP in Firefox, even in a very simple form. As it turns out, we had much bigger plans and that's what I'm sharing here today.

Over the coming weeks, Mozilla will begin trialing integration between HIBP and Firefox to make breach data searchable via a new tool called "Firefox Monitor".

Here's what it looks like:



This is major because Firefox has an install base of hundreds of millions of people which *significantly* expands the audience that can be reached once this

feature rolls out to the mainstream. You can read <u>Mozilla's announcement of the new feature</u> and how they plan to conduct the testing and rollout.

I'm really happy to see Firefox integrating with HIBP in this fashion, not just to get it in front of as many people as possible, but because I have a great deal of respect for their contributions to the technology community. In particular, Mozilla was instrumental in the birth of Let's Encrypt, the free and open certificate authority that's massively increased the adoption of HTTPS on the web. Arguably, the work done by Mozilla's Josh Aas and Eric Rescorla (still the Mozilla CTO today) has been one of the greatest contributions to online privacy and security we've seen and Mozilla remains a platinum sponsor to this day. They've also been instrumental in helping define the model which HIBP uses to feed them data without Mozilla disclosing the email addresses being searched for. I'm going to talk more about the mechanics of that model in a moment but first, let me talk about 1 Password:

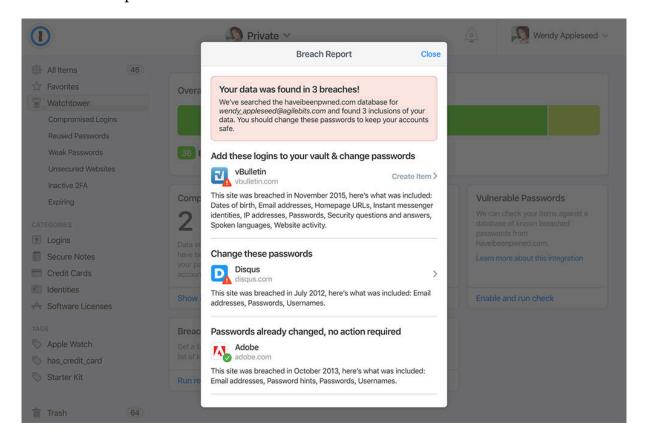
1Password

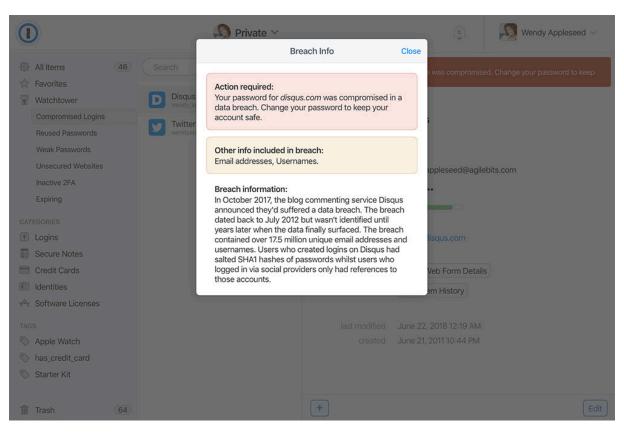
My relationship with 1Password stretches all the way back to 2011 when I came to the realisation that the only secure password is the one you can't remember. Over the last 7 years, I've continued to buy their product and use it every single day across all my devices and my entire family's devices. In February, only the day after I launched Pwned Passwords V2, 1Password turned around and built it into their product so that users of the password manager could see if their password had been previously exposed in a breach. That effort was a large factor in my choosing 1Password to partner with HIBP back in March and since that time, they've built Pwned Passwords into the desktop apps for Mac and Windows and provided the ability to check all your passwords in one single go. But today, we're announcing something much bigger:

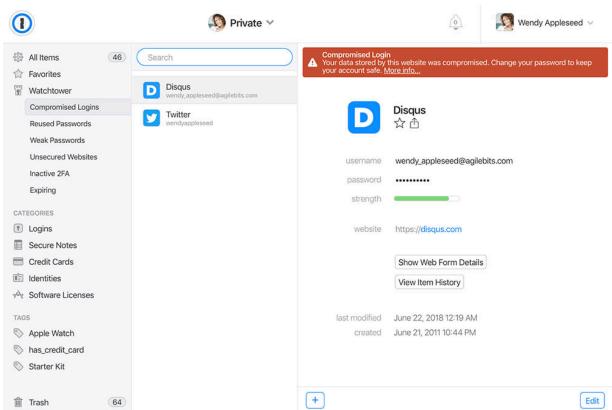
As of now, you can search HIBP from directly within 1Password via the

Watchtower feature in the web version of the product.

This helps Watchtower become "mission control" for accounts and introduces the "Breach Report" feature:







As with Pwned Passwords, by pushing this out in the web-based version of the product they can get it to customers quickly then over time, bake it right into the desktop versions as well. There's also a bunch of other ways 1Password can use the data to streamline how users protect their accounts and that's something we're actively discussing. I expect we'll see the existing functionality enhanced in the not too distant future.

If you're a 1Password user you can use this feature right now, just head on over to the 1Password login page. And if you're not already putting all your passwords in 1Password, go and grab a free trial and give it a go. You can also find a more detailed write-up on 1Password's implementation in the very aptly titled blog post: we shall fight on the breaches (why didn't I ever think of that?!)

Enabling Anonymous Searches with k-Anonymity

I want to talk about protecting the identities of Firefox and 1Password users because more than ever - and regardless of where you are in the world - we're becoming increasingly conscious of our online privacy. We're also becoming increasingly connected and sharing unprecedented volumes of data which, let's face it, isn't exactly analogous with privacy and anonymity! But we *can* have both and I want to illustrate that by talking about <u>the Pwned Passwords model</u> for a moment.

When this feature launched, Cloudflare (hat-tip again to <u>Junade Ali</u> there) did some great work on a "k-anonymity" model which works like this: when searching HIBP for a password, the client SHA-1 hashes it then takes the first 5 characters and sends this to the API. In response, a collection of hashes is returned that match that prefix (477 on average). By looking at the hash prefix sent to the service, I have no idea what the password is. It *could* be any one of those 477 or it could be something totally different, I don't know. Of course, I

could always speculate based on the prevalence of each password but it would never be anything more than that - speculation. (Just to add to that, I've never got any idea of the username attached to the password either so even if I take an educated guess at it, there's nothing I can actually do with it.)

The email address being searched for by Firefox and 1Password works in the same fashion, albeit it with slightly different numbers due to the significantly larger data set at play. When I processed the source HIBP data in preparation for this feature, out of the 5B records in the system at the time there were 3.1B unique email addresses. (In other words, each address has been in an average of 1.6 data breaches.) I took each one of those 3.1B addresses, hashed it and stored it in a new data construct I'll talk about later. That gave me a repository to search against, now let's cover the mechanics of that search:

For the purposes of anonymity, I needed to decide on how many characters of the SHA-1 hash to allow searching by such that a sufficiently large number was returned to have no reasonable way of knowing which address was searched for, but also for the system to respond quickly. For Pwned Passwords, that number was 5 characters resulting in 16 ^ 5 possible search ranges which, across a data set of 500M records, meant the aforementioned 477 results per range. However, if I'd used 5 chars with the 3.1B email addresses, each range would contain an average of almost 3K results which is starting to get pretty sizeable.

Ultimately, I settled on 6 characters which means 16 ^ 6 possible ranges with an average of 185 results per range. Now, on the one hand you might say "that's less than Pwned Passwords therefore provides less protection", but it's a bit more nuanced than that. Firstly, because this number will grow *significantly* over time; more data breaches means more new email addresses means larger results in the range search. More importantly though, email addresses are *far* less predictable than passwords; as I mentioned earlier, if I was to spy on searches for Pwned Passwords (and I don't, but this is the threat k-anonymity is protecting us from), the prevalence of passwords in the system beginning with that hash can indicate the *likelihood* of what was searched by. But when we're

talking about email addresses, there's no such indicator, certainly the number of breaches each has been exposed in divulges nothing in terms of which one is likely being searched for.

Here's what a search for an email address ultimately looks like:

```
Address: test@example.com

SHA-1 hash: 567159D622FFBB50B11B0EFD307BE358624A26EE

6 char prefix: 567159

API endpoint: https://[host]/[path]/567159
```

Which results in a response containing the following:

```
- {
      hashSuffix: "D42AE
    - websites: [
          "Tumblr",
          "WeHeartIt"
      ]
  },
- {
      hashSuffix: "D5E66
    - websites: [
          "OnlinerSpambot"
      ]
  },
  {
      hashSuffix: "D622FFBB50B11B0EFD307BE358624A26EE",
    - websites: [
          "000webhost",
          "2844Breaches",
          "7k7k",
          "8tracks",
          "Adobe",
          "AntiPublic",
```

In this case, the searched address is the last one because the hash suffix matches with the SHA-1 hash of <u>test@example.com</u>. The associated websites next to that hash are the ones that the email address appeared in and can be matched back to the full breach details via <u>the public breach list API</u>.

Unlike the way 1Password implements Pwned Passwords by calling the API directly from the client, this model is only consumable by an authenticated request from Firefox's or 1Password's infrastructure. What this means is that consumers only ever call their API then they call HIBP's API. Both these organisations implement all the same controls that HIBP does on the existing public email search API, namely the rate limit, not returning sensitive breaches to non-verified addresses and employing a range of other abuse-protection mechanisms. Why not call the API directly from each client? Let's talk about that next:

The Viability of Using This API Publicly

There's one fundamentally important (and perhaps quite obvious) reason why I don't expect to make this service available publicly: it could *massively* accelerate enumeration activities. Back in 2016, I implemented a rate limit on the public API to greatly reduce the potential to abuse the service. This meant the ability to check records was limited to 1 request every 1,500ms. If I was to offer the k-anonymity service publicly, that jumps massively to 185 every 1,500ms (and it will grow as the data does) because that's how many results are returned. In fairness, you'd only get back hashed suffixes of email addresses but if someone had a *massive* list of them they wanted to work through (and that's one of the key patterns the rate limit is designed to curtail), they could hash the lot then grab those first 6 chars of each and get back a bunch of results in one go. Granted, they almost certainly wouldn't have the source email addresses of all 185 returned suffixes, but it still provides a vector to greatly accelerate search rates.

A (somewhat) middle ground, however, would be to allow searching the repository of hashed addresses by complete hash. Rather than the current model which needs a full email address, now that I have a mirrored data set containing only hashes I could always roll the search over to query that repository instead. I could also adapt the web front end to hash an entered email address client-side then only send that to the server. I actually closed a User Voice idea along these lines (and you can read about why in that idea), but now that I have that collection of hashes already in storage, it'd be trivial to stand up an endpoint to query it even if there's not always a lot of privacy upside. The main reason I can't do that immediately is that the email addresses for the paste service are not hashed and the existing search model on the website needs to search both sources.

And just on that, let me refer quickly back to my post on <u>The Ethics of Running</u> a <u>Data Breach Search Service</u>. This explains in detail *why* the service allows addresses to be searched in the way it presently does and provides both technical and logical reasons. Do read that if you're curious, it was very carefully thought-out and explains the detailed thinking behind it.

Scaling Searches with Azure and Cloudflare

As with the k-anonymity model itself, I've leaned heavily on the Pwned Passwords experience in designing how this model works. <u>Azure Functions</u> provide the API layer (they're serverless and scale beautifully) and Cloudflare does the reverse-proxying and caching. Unfortunately, I can't cache anywhere *near* as aggressively as with Pwned Passwords because instead of 16 ^ 5 different ranges (and therefore unique request patterns), it's 16 ^ 6. Also, I have a one-month cache expiration on Pwned Passwords because they rarely change but at present, only a 15-minute cache expiration on the email address search. I didn't

want someone searching for a breach I'd just loaded, finding it in HIBP then *not* finding it in Firefox or 1Password.

One major difference between Pwned Passwords and this feature here is the underlying storage construct. For the passwords, I ended up putting them all in blob storage so there's one file per hash range. That was faster than the original Table Storage construct I used and because the files rarely ever change (I'll probably only update the passwords a couple of times a year), they could remain relatively static files. For emails in breaches, however, I don't have that luxury; breaches are continually loaded into the system and I needed a queryable construct that allowed for fast inserts and reads which means that like the existing HIBP email address storage, the hash ranges for this service are also in Azure Table Storage. It's an entirely separate table to the one that's been there from day one and just holds email addresses so obviously that's also impacted the data load process which now needs to do twice as much record handling.

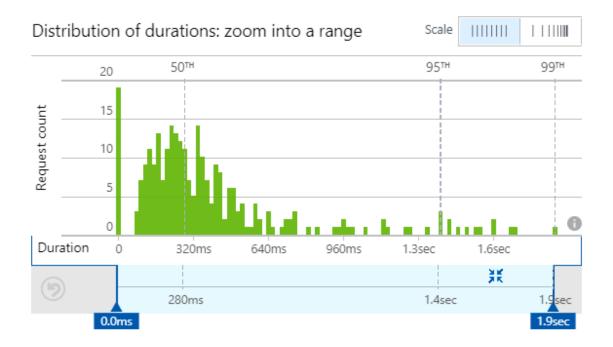
What this means is that the entire record for <u>test@example.com</u> looks like this:

Partition key: 567159

Row key: D622FFBB50B11B0EFD307BE358624A26EE

Websites: [delimited array of pwned site names]

Partitioning in this fashion means that when a search is done, it's easy to return every single suffix for the hash prefix which is used as the partition key. Some quick performance testing last week resulted in the following:



That's fine as a starting point and median load time of 280ms means that even with network latency and processing on Firefox's and 1Password's end should mean results are turned around in well under a second. Functions also seem to accelerate in execution time as infrastructure warms up so I expect we'll only see improvements in this area. I'll share some perf info once the volumes really ramp up.

Finally, like the existing search services and regardless of the fact requests only ever contain 6 characters of a hash, no searches are ever explicitly logged. They'll pass through very short-term transient logs and that's it - all I'll have is very broad-brush stats on things like the number of calls and the duration of executions. So, in summary, everything will be fast, efficient and above all, anonymous.

The Notification Service

Let me talk briefly about one last thing that's on the horizon - notifications. The service above is a point-in-time representation of breach state which is great,

but clearly, it's a state that will continue to evolve over time. There needs to be a construct to notify Firefox and 1Password of *deltas* to the data set as new breaches are loaded so that in the future, they have the ability to offer subscriptions in the same way as the HIBP notification service. Re-querying all the data every time a breach is loaded would be *massively* inefficient so that's not going to happen, we need to be smarter than that.

A few years back <u>I wrote about a callback model</u> and an equivalent paradigm will be used here. What this boils down to is the ability to subscribe a hash prefix for notifications so in that <u>test@example.com</u> scenario, "567159" would be subscribed then a callback sent to Firefox or 1Password when an address that hashes down to that prefix is loaded. The callback will contain the name of the breach that's just been loaded and all the hash suffixes for that prefix that were found in it.

I'm building that feature out next and in time I expect we'll see that flow through to Firefox and 1Password. It's a neat way of ensuring that anonymity is still protected *and* that subscribers of those services stay abreast of security incidents that impact them.

Summary

As HIBP grows, I keep coming back to this question:

How can HIBP do good things for people in the wake of bad events?

I'm really happy this initiative furthers that objective and does it in a way that puts privacy first. The leverage these two organisations have to drive positive outcomes in the wake of data breaches is massive and I'm *enormously* excited to see the impact they both make in partnership with HIBP.

Comments

This is *incredibly* exciting! Baking these types of security and privacy services directly into browsers will exponentially increase the chance they end up helping the average web user who doesn't know they even exist. I hope this trend continues and browsers leverage services like HIBP to help users become more secure online. For example, I would love to see FireFox (and the other browsers!) use the Pwned Passwords API to natively do something like the awesome PassProtect web extension (https://www.passprotect.io/), which informs users if they are about to create a password that has previously been compromised in a data breach.

Troy, what types of features/functionality like this do you think make sense to bake into browsers moving forward?

Troy: It's a good question, browsers are progressively building in more and more account management functionality so things like Pwned Passwords and what Firefox is doing here make a lot of sense. In essence, anything that saves people from themselves (i.e. helps them make wise security choices) is a good thing IMHO.

I'm appalled about this project. Basically, it tries to help people unable or unwilling to chose proper passwords and to properly administer them all. If every on-line service was obliged to duplicate its database (just in case of being hacked), the service could inform the account/mailbox/password owner by contacting the owner through a second mail address, for which the service has no password, but the address itself, only. At least this would do the same as this new firefox integration project can do (or will do). But now, inviting hundreds of millions uf irresponsible users to somehow publish their mail addresses and passwords in advance, must lay the base of future data breach of astronomical dimension, you bet. For me, I don't need such crap, don't need password managers of any kind, and I would not come to the idea of letting third parties of any kind administer or store my many, countless passwords and accounts, since this is so easy done by myself - and way more reliable and secure, too. It's hard to believe that there is a market out there for mother and nanny

credulous, naive people in such proportions.

Troy: I'm not sure precisely what it is you disagree with - nothing identifiable is sent to HIBP, is it password managers themselves you're displeased with? They're the best option we have by a long shot when it comes to passwords, certainly your brain is a very poor storage mechanism that always results in serious compromises to password strength.

I'm safeguarding my passwords and all sorts of countless on-line access credentials within an encrypted file container (protected by a long and complex story-like password string no library would ever hold), containing a simple editor file with all these countless login data, passwords, on-line credentials altogether in plaintext. This container is always separated, off-line, never located on a running machine/device with internet access. Furthermore, this encrypted file is backed-up on many storage devices, at different places, continuously maintained, which is easily done, and which is way more secure than what this firefox project (or other dedicated password managers, or cloud-based solutions) can offer. I'm safe and comfortable with this security policy for about 20 years, while I have seen many others losing all their critical data, by airily and careless data handling, all of sudden, repeatedly.

What I'm criticising is two things: This Firefox project promotes a dangerous understanding of password and privacy security that can't be guaranteed or sustained (by Firefox and any other on-line, third party, cloud service), while it fosters sheer lazyness and carelessness at the side of so many, privacy-unaware on-line Users. Instead, Users should care for their privacy and password security on their own. I'm afraid, at the end of the day, this fatal "pairing" (illusionary security promises versus Users' credulity and carelessness) aims silently at what is going on in the Web since long ago, the longer the more, which is profane, massive User data collection, being prone to abuse (or say monetarising) anytime, if a cloud-based "service provider" decides to do so. Secondly, this project drives the risks of massive data breach just to another, much greater level, since, at this point, hundreds of millions Users will be involved eventually, and since, empirically, it is not in question whether such an incident will happen, but only, when it will happen. The meaning of both objection points is

what let's me shiver.

Troy: So it's just a "roll your own" password manager then and as you say, you take care of all the crypto, backups, syncing etc. That's fine for someone savvy enough, I'm sure you can see why that's an exceptional situation though.

I can't speak for Firefox's implementation, but certainly 1Password's is extremely well thought out, independently audited and extensively documented. Even a total compromise of the 1Password service wouldn't expose user passwords due to the way they're stored. This is a great whitepaper on that: https://lpassword.com/files...

I'm not sure which incidents you're referring to when you say "many others losing all their critical data", which mainstream password manager has this occurred to? But hey, if you can find a way to do that with 1Password then you should definitely go and pick up the \$100k bounty on offer! https://www.digitaltrends.c...

As much as I feel that password security is absolutely needed, I feel like this is the wrong approach to get there. We're basically telling users to trust pushing your hashed password through an interface to a third party over a network. A localized, updatable dictionary of bad passwords I would be totally cool with, but this is giving up security to maintain security. It's a contradiction in policies.

No security person would ever agree to this being done as a home grown solution in house, which is why it shouldn't be done as a solution for anyone. The minute any part of this gets hacked or compromised, criminals will be able to access at least who the weakest users/orgs are as far as infosec is concerned. Although password lists are out there, those are at least dated. I want to believe that all you'd have to do is get a man in the middle point, get the size of the organization and their password change policies, then look at how large their traffic to pown password checks have been done to see if they have particularly vulnerable users. They didn't pass on the first, second, third, fourth checks? That's a lot of round trip traffic that

could be monitored.

Maybe what you're really doing by this is sharing with companies internally who has good infosec and not. While I don't think having sloppy employees is a particularly good thing, I wouldn't want to be the one to play snitch for corporate either.

Troy: You really don't want half a billion passwords sitting locally on every consumer of this service which is why the API is getting so much traction. And others have said, we're only talking about hash prefixes which by the design of k-anonymity, don't give you anything usable even before you consider the fact you can't explicitly tie the password back to a user anyway.

Pwned Passwords is now pushing 8M+ requests a day to the API and is used by a growing number of large and small orgs alike. It has traction precisely because it provides anonymity; plus being fast and free helps a lot too

It's only the first five characters of the SHA-1 hash that's being sent. Why is this problematic?

It's problematic being sent over a network. Are traffic packets totally anonymous too? You can automatically see who is using it and how often. Besides that, it's not rocket science to make a lookup table for the first five characters of the SHA-1, which then can link to all known possibilities for weak passwords. None of this screams secure - just probably more secure than using passwords from an insecure password dump list. You shouldn't be transmitting any hint of something password related you don't need to.

Troy: This model isn't sending anything related to passwords, it's about email addresses. Or are you referring to the Pwned Passwords service instead? I'm trying to understand which service you're referring to and what real-world risk you're seeing.

I still feel wary of the k-Anonymity security model. I trust Troy Hunt, but I wouldn't give anyone my passwords. Now, what I worry about is if someone got a hold of the web logs from Azure or something, just say. And they can correlate via IP address those API requests back to you, and say your password is compromised, now they know your password is one of 400... It just feels dirty sharing even part of a SHA1 hash of your passwords...

Troy: But they don't know your password is "one of 400", they'd know the first 5 chars which means they know the *hash* is one of 16 ^ 35 different ones (SHA-1 hashes are 40 heaxadecimal bytes), which, for the sake of completeness, means they can narrow it down to one of

This is a grey area which I honestly wouldn't approach. The possibility of a service like this being directly or indirectly hijacked is much too high. You're grooming a pool of users to believe it's okay to send passwords off to a non-authentication server, which is also bad. Even the metadata from a password should not be sent to additional machines or over additional networks. It's a poor and sloppy practice.

Normally, when a company loses personal information the liability is quite limited. However, when it's proven that many complex passwords have only been shared directly to the authentication service, and to a password check site (such as this), it will be very different. Liability will be high.

Troy: Let's try it like this: what, precisely, is the worst possible outcome you can conceive of if HIBP is hijacked, as you put it, and someone begins obtaining the first 6 chars of SHA-1 hashes of email addresses? Or is it more that you're against password managers themselves?

Epilogue

All of this was very fortunate timing; it was only 4 months earlier that I'd launched the Pwned Passwords k-anonymity search so there was a very simple, very easily implemented privacy model ready to go. Despite dedicating a section of the blog post to why I couldn't do this, I'd later get requests to make it available publicly as an alternative to the existing API that allowed searching by email address alone. The challenge with this remained that because each search returns hundreds of results, I couldn't roll it out en masse to the public as it would completely undermine the premise of rate-limiting. The whole idea of the rate limit was to chop out a huge portion of the abusive patterns I was observing and whilst k-anonymity did wonderful things for privacy, it could also supercharge abuse.

On a whim, I just went and looked at the performance of the Azure Function supporting this API. It must have been 2 years since I'd done this, I just didn't have a need before now. Turns out that median load time is down from 280ms in the blog post to less than half that at 130ms over 30 days of queries encompassing over 6 million requests. This is one of the things I absolutely love about the chosen architecture of this entire project - I never even think about it anymore. I pay the bills and stuff just works

As I write this, I can't help but come back to the same premise that I've reiterated many times over in this book: activities and relationships that I established purely through the love of technology and a desire to make a

positive difference eventually turned into something much more valuable. I don't just mean in a monetary sense (both Mozilla and 1Password do pay for the services they use), but in the sense that they made a difference to the world. How many people have learned of a data breach through one of these services? Millions? Tens of millions? I'll never know exactly; I just know it's a big number and that's awesome.

THE EFFECTIVENESS OF PUBLICLY SHAMING BAD SECURITY

I started writing this blog post angry. It wasn't that I was pissed at the way organisations sometimes approached and then publicly communicated security, rather I was pissed that after I'd publicly call them to account, I'd occasionally be criticised for doing so. It just felt like this overwhelming modern-day self-righteousness where no matter how badly someone screws up, there's always a brigade of people standing by asking you to be nice to them and not hurt their feelings. I was fed up.

I wanted to write this post so that every time someone piped up and said, "you shouldn't publicly shame companies for doing stupid security things after they do stupid security things", I'd have a post explaining why they were wrong. A thoughtful, well-articulated post that drew on previous examples illustrating my point. And there are many examples of where this has worked beautifully. What this post helped me do is to shift the discussion from "does shaming work" to "is shaming nice" and frankly, I really don't care about the latter if this approach achieves the former. And it does.

11 SEPTEMBER 2018

ere's how it normally plays out: It all begins when a company pops up online and makes some sort of ludicrous statement related to their security posture, often as part of a discussion on a public social media platform such as Twitter. Shortly thereafter, the masses descend on said organisation and express their outrage at the stated position. Where it gets interesting (and this is the whole point of the post), is when another group of

folks pop up and accuse the outraged group of doing a bit of this:

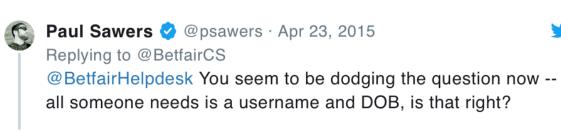


Shaming. Or chastising, putting them in their place or taking them down a peg or two. Whatever synonym you choose, the underlying criticism is that the outraged group is wrong for expressing their outrage towards the organisation involved, especially if it's ever construed as being targeted towards whichever individual happens to be the mouthpiece of the organisation at the time. Shame, those opposed to it will say, is not the way. I disagree and I want to explain - and demonstrate - precisely why.

Let's start with a few classic examples of the sort of behaviour I'm talking about in terms of those ludicrous statements:









@psawers Yes, but they would need to attain this information through you, which once again, is a breach of our terms.

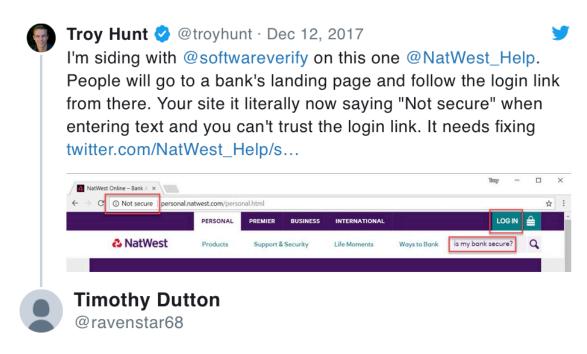
○ 6 7:44 AM - Apr 23, 2015

See the theme? Crazy statements made by representatives of the companies involved. The last one from Betfair is a great example and the entire thread is worth a read. What it boiled down to was the account arguing with a journalist (pro tip: avoid arguing being a dick to those in a position to write publicly about you!) that no, you didn't just need a username and birth date to reset the account password. Eventually, it got to the point where Betfair advised that providing this information to someone else would be a breach of their terms. Now, keeping in mind that the username is your email address and that many among us like cake and presents and other birthday celebratory patterns, it's reasonable to say that this was a ludicrous statement. Further, I propose that this is a perfect case where shaming is not only due, but necessary. So I wrote a blog post..

Shortly after that blog post, three things happened and the first was that it got press. The Register wrote about it. Venture Beat wrote about it. Many other discussions were held in the public forum with all concluding the same thing: this process sucked. Secondly, it got fixed. No longer was a mere email address and birthday sufficient to reset the account, you actually had to demonstrate that you controlled the email address! And finally, something else happened that convinced me of the value of shaming in this fashion:

A couple of months later, I delivered the opening keynote at OWASP's AppSec conference in Amsterdam. After the talk, a bunch of people came up to say g'day and many other nice things. And then, after the crowd died down, a bloke came up and handed me his card - "Betfair Security". Ah shit. But the hesitation quickly passed as he proceeded to *thank me* for the coverage. You see, they knew this process sucked - any reasonable person with half an idea about security did - but the internal security team alone telling management this was not cool wasn't enough to drive change. Negative media coverage, however, is something management actually listens to. *Exactly* the same scenario played out at a very similar time when I wrote about how you really don't want bank grade security with one of the financial institutions on that list rapidly fixing their shortcomings after that blog post. A little while later at another conference, the same discussion I'd had in Amsterdam played out: "we knew our SSL config was bad, we just couldn't get the leadership support to fix it until we were publicly shamed".

I wanted to set that context because it helps answer questions such as this one:



Why does it often take being named and shamed before they actually do something about these vulnerabilities. Still nice to see they have actually changed the site now.

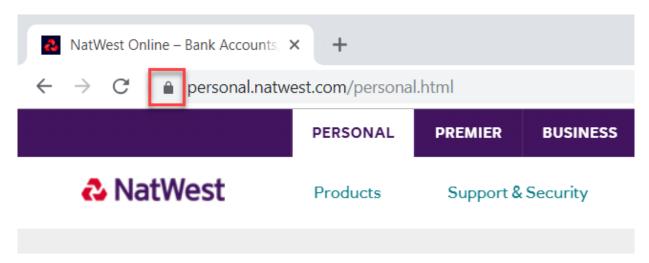
○ 3 5:02 AM - Dec 17, 2017

What public shaming does is appeals to a different set of priorities; if, for example, I was to privately email NatWest about their lack of HTTPS then I'd likely get back a response along the lines of "we take security seriously" and my feedback would go into a queue somewhere. As it was, the feedback I was providing was clearly falling on deaf ears:



And now we have another perfect example of precisely the sort of response that *needs* to be shamed so <u>NatWest earned themselves a blog post</u>. How this changed their priorities was to land the negative press on the desk of an executive somewhere who decided this wasn't a good look. As a result, their view on the security of this page is rather different than it was just 9 months ago:

○ 5 9:59 AM - Dec 12, 2017



Now I don't know how much of this change was due to my public shaming of

0

their security posture, *maybe* they were going to get their act together afterward anyway. Who knows. However, what I *do* know for sure is that I got this DM from someone not long after that post got media attention (reproduced with their permission):

Hi Troy, I just want to say thanks for your blog post on the Natwest HTTPS issue you found that the BBC picked up on. I head up the SEO team at a Media agency for a different bank and was hitting my head against a wall trying to communicate this exact thing to them after they too had a non secure public site separate from their online banking. The quote the BBC must have asked from them prompted the change to happen overnight, something their WebDev team assured me would cost hundreds of thousands of pounds and at least a year to implement! I was hitting my head against the desk for 6 months before that so a virtual handshake of thanks from my behalf! Thanks!

Let me change gear a little and tackle a common complaint about shaming in this fashion and I'll begin with this tweet:





Ok England, look, this sort of stuff was funny for a while and I appreciate the laughs, but it's starting to get a bit ridiculous. Can one of you please pop down to @santanderukhelp HQ and straighten this mess out?

Santander UK Help 🤣 @santanderukhelp

Replying to @markhood

Hi Mark, we would never recommend using 3rd party password managers. It is no longer possible to use these for security reasons. ^JM

0

Notwithstanding my civic duty as an Aussie to take the piss out of the English, clearly this was a ridiculous statement for Santander to make. Third party password managers are *precisely* what we need to address the scourge of account takeover attacks driven by sloppy password management on behalf of individuals. Yet somehow, Santander had *deliberately* designed their system to block the ability to use them. Their customer service rep then echoed this position which subsequently led to the tweet above. That tweet, then led to this one:





Replying to @troyhunt @santanderukhelp

Please, just not another witch hunt on some poor clueless Customer Service rep... :(

Andy is concerned that shaming in this fashion targets the individual behind the social media account (JM) rather than the organisation itself. I saw similar sentiments expressed after T-Mobile in Austria defended storing passwords in plain text with this absolute clanger:

<u>@Korni22</u> What if this doesn't happen because our security is amazingly good? ^Käthe

-- T-Mobile Austria (@tmobileat) April 6, 2018

In each incident, the respective corporate Twitter accounts got a lot of pretty candid feedback. And they deserved it - here's why:

These accounts are, by design, the public face of the respective organisations. Santander literally has the word "help" in the account name and T-Mobile's account indicates that Käthe is a member of the service team. They are absolutely, positively the coal faces of the organisation and it's perfectly reasonable to expect that feedback about their respective businesses should go to them.





Replying to @markhood and 2 others

Social media accounts are the public face of an organisation. Their specific remit is to engage with the public who'll likely have something to say about this policy.

○ 7 3:15 AM - Apr 18, 2018 · Melbourne, Victoria

0

This is not to say that the feedback should be rude or abusive; it shouldn't and at least in the discussions I've been involved in, that's extremely rare to see. But to suggest that one shouldn't engage with the individuals controlling the corporate social media account in this fashion is ludicrous - that's *exactly* who you should be engaging with!

A huge factor in how these discussions play out is how the organisations involved deal with shaming of the likes mentioned above. Many years ago now I wrote about how customer care people should deal with technical queries and I broke it down into 5 simple points:

- 1. Never get drawn into technical debates
- 2. Never allow public debate to escalate
- 3. Always take potentially volatile discussions off the public timeline
- 4. Make technical people available (privately)
- 5. Never be dismissive

Let me give you a perfect example of how to respond well to public shaming and we'll start with my own tweet:





What is it with the anti-password-pasters today?! How is this sentiment permeating into organisations like @medibank in an era of so many password abuses?

Hi Ben, we ask for the password to be typed rather than pasted to make sure that it's entered correctly. I understand that pasting would be easier – we just want to make sure that additional characters/spacing are not pasted too. Thanks for the feedback and writing to us. -Ane

Business as usual there, just another day on the internet. But watch how Medibank then deals with that tweet:





A

Replying to @troyhunt

Hi Troy, We just wanted to let you know that we've checked in with our digital team and they've let us know that they are already in the process of resolving this. We'll be deploying the ability to paste in about two weeks. Thanks again for the feedback! Kindly, Sam.

0

And in case you're wondering, yes, <u>I did give them an e-pat on the back for that</u> because they well and truly deserved it! The point is that shaming, when done

right, leads to positive change without needing to be offensive or upsetting to the folks controlling the social accounts.

The final catalyst for finishing this blog post (I've been dropping examples into it since Xmas!) was a discussion just last week which, once again, highlighted everything said here. As per usual, it starts with a ridiculous statement on security posture:



MARK GOOK = Mark Williams-Cook ... · Sep 4, 2018 > I wrote a post asking @tvlicensing to please secure (https) their website when requesting customer names, emails addresses and bank details - it's not! i83.co.uk/why-tvlicensin... #infosec



Why tvlicensing.co.uk are processing millions of cust...
Despite the claims on their website, tvlicensing.co.uk are sending personally identifiable information and bank details i83.co.uk



Our website is secure and security certificates are up to date. Pages where customers enter data are HTTPS. Non HTTPS pages are safe to use despite messages from some browsers (e.g. Chrome) that say they are not.

Shaming ensues (I mentioned my Aussie civic duty, right?!):





I don't get British humour

TV Licensing @tvlicensing Replying to @thetafferboy

Our website is secure and security certificates are up to date. Pages where customers enter data are HTTPS. Non HTTPS pages are safe to use despite messages from some browsers (e.g. Chrome) that say they are not.

0

Once again, the press picks it up and also once again, people get uppity about it:





Replying to @CountVice and 2 others

Also this is a social media account not a first response security account. Yes they are wrong, but as with T mobile and others- are we using a social media mgr to shame an org? Yes we need better awareness. But shame isn't the way.

○ 1 2:44 PM - Sep 5, 2018

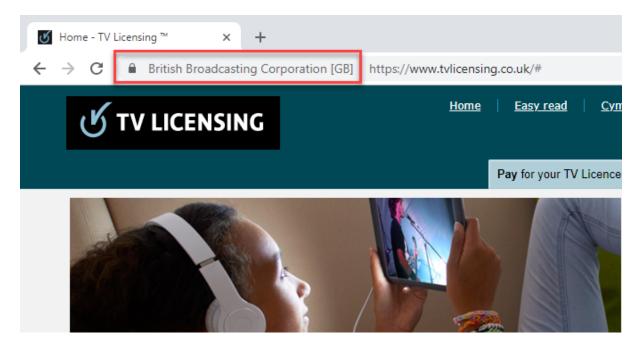
0

'these guys' = some person working a minimum wage customer service job + raising the issue led to the issue being resolved. Calling them 'not bright' when they have to deal with whatever questions get thrown their way despite no real investment in them is not nice.

-- Chris (@Modularized) September 9, 2018

And just to be clear, stating that "Non HTTPS pages are safe to use despite messages from some browsers" is not a very bright position to take whether

you're on minimum wage or you're the CEO. Income doesn't factor when you make public statements as a company representative. Predictably, just as with all the previous example, positive change followed:



That whole incident actually <u>turned out to be much more serious than they originally thought</u> and once again, the issue was brought to the forefront by shaming. I've seen this play out *so* many times before that frankly, I've little patience for those decrying shaming in this fashion because it might hurt the feelings of the very people charged with receiving feedback from the public. If a company is going to take a position on security either in the way they choose to build their services or by what their representatives state on the public record, they can damn well be held accountable for it:



y

Replying to @troyhunt and 3 others

I'm *absolutely* fed up of social media managers/comms teams taking control and making erroneous statements. If they have the balls to say something that's demonstrably false and won't back down when shown proof, be it on their head.

0

Whether those rejecting shaming of the likes I've shared above agree with the practice or not, they can't argue with the outcome. I'm sure there'll be those that apply motherhood statements such as "the end doesn't justify the means", but that would imply that the means is detrimental in some way which it simply isn't. Keep it polite, use shaming constructively to leverage social pressure and we're all better off for it.

Comments

Note the difference in response depending on who does the shaming - when someone unknown points out a flaw, nothing is done. But when someone with a significant public following points it out the issue gets resolved. This makes public shaming only effective if you are famous.

Troy: I think a more accurate way of describing that is the more exposure the shaming gets, the more pressure there is on the organisation to reform. Many of these incidents begin by someone with less influence than me raising the issue then myself (or others) circulating it further. Imagine it from the company's

perspective - one person complains and it's easily dismissed, but get hundreds or thousands of tweets about it and it changes the priorities pretty quickly!

Of course - but you and others with your level of following can't have eyes everywhere. I'm sure many reported issues fall through the cracks, and the less known people don't know how/ have the motivation to try to increase publicity.

Troy: Of course.

While that is true, it only takes including <u>Troy Hunt</u> to the twitter to get his attention and if he agrees he can then choose to lend his weight to the "education" campaign.

Troy: Totally! So long as it's not just another "this company limits passwords to n chars long" issue! Seriously, it's the egregiously bad or stupid stuff that gets traction.

One note regarding password managers.

When your users are mostly geeks and academics then you can implement any advanced solution and give sensible advices like "use password managers and generate passwords with at least 200 bits of entropy".

When your users represent full range of education levels and computer literacy (ie. anyone) - then you have to do the analysis:

on one side you have "no official support for password managers" (it means large number of password reuse and passwords leaks from other sources + which you can mitigate with 2FA), on the other side you have "advice to use password managers" (it means that users will download malware impersonating password manager, their passwords will leak from low quality managers, they will upload their vaults to the cloud or synchronize with public

computers, they will reuse their favourite password as the master password or they will forget their master passwords and expect you to sort it out because it was your advice).

Having above in mind - for many users using password manager is much less secure than writing the password on paper and keeping it in drawer.

Troy: If people were writing unique passwords on paper and putting it in drawers we'd be *way* better off than we currently are! But on the whole they're not.

Every other issue you've raised can be managed if there's a desire to adapt behaviour. That's still a barrier, of course, but there's a lot of headway being made. For example, increasingly when I go into organisations they're rolling out formally approved password managers as part of their SOE. Further, companies like 1Password are providing free family licenses alongside corporate subscriptions to help change user behaviour in the home too. And as for downloading malicious software, that's increasingly hard in a "mobile first" era when apps are so frequently taken from official stores.

As I've said before, password managers don't have to be perfect, they just have to be better than not having one: https://www.troyhunt.com/pa...

Using (or even forcing) managers within organisations is not a problem, it's well known path, it's secure environment, automated deployment etc. But it's much different for clients. Nice example is 2 years old attack - after some banks made a deal with well known antimalware company to promote their product among clients we've seen really large wave of different attacks aimed at clients (including pc malware and massive mail phishing campaigns - like "you've heard from your bank about this new shiny antivirus - so that's it, this is the link, click and follow instructions"). After that many organisations will think twice before advertising installation any security-related product.

https://www.cert.pl/en/news...

Troy: I think the point you're making is that if people improve their security posture with a password manager then attackers will begin targeting them. That's always the way though; for every security measure we implement, someone tries to circumvent it. But it greatly raises the bar and certainly as of now, they're still rare enough that those without them will remain the weakest links and the most easily targeted.

There's a problem not addressed here: when someone walks up to you and says "Thanks for the public outcry, you managed to convince the management" that means said company had an inadequate management with a total lack of technical understanding. Shaming is nice, but we didn't solve the problem that those people are charlatans and most likely the next issue would suffer from the same. In short they should be fired.

Troy: Depends on how you look at it: often the underlying issue is that there are bigger priorities and the shaming readjusts those priorities. But I do agree it's often also reflective of shortcomings on behalf of the the organisation's security posture, although in one of these cases that ultimately led to me doing a couple of workshops (at their request), so shaming can also drive positive change that way too.

To those who decide to shame others, beware: Some companies will fight back. It's a stupid move in the long run for them, but in the short run, you could find yourself having to hire an attorney to defend yourself, even though you may feel you have done nothing wrong.

To those who delight in being that kid who pointed out the nudity of the emperor, just remember this: Make sure it is a very wide audience that cares. Make sure that you have documented and exhausted all other means of notifying the offender. Make sure that you have safe legal ground to stand on.

I say this as one who has been threatened by lawyers on behalf of offending companies.

And by the way, to you low life attorneys who perpetrate those "cease and desist" letters: You are the reason why people view attorneys as people of such ill repute. This is why your profession is considered one of the bottom feeders of society. The next time some chump client comes along asking you to write those letters, know in the long run, you will lose far more than any fees you may collect.

Troy: Let's be really clear about this: all the examples I've shared are either of publicly observable security posture or in response to statements made by the organisation. None of this trumps responsible / coordinated disclosure and none of this condones defamation, as someone else suggested.

In these circumstances, I disagree with "exhausting all other means". Let's take this from only an hour ago as an example:



I feel in no way whatsoever compelled to invest effort in tracking down the people within the organisation who are responsible for this and appealing to them. The Commbank have very clearly stated their position in the public domain and I'm calling attention to that. That's a perfectly reasonable position to take and the onus should be on *them* to rectify it, not on me to do anything "exhausting".

Epilogue

It was around this time in my blogging life that I started embedded a lot more tweets into blog posts. Obviously, there's a bunch of company tweets (namely the ones I used as examples that should be shamed), but there are also tweets from individuals. Sometimes (as is the case in this blog post), those tweets expressed views I disagreed with, and I'd use the blog to highlight why. I've occasionally received criticism for this as some deem it as shaming the individual (there's a theme emerging here...), whether that be in this blog post or subsequent ones following the same approach. But I vehemently disagree with the premise that someone's public tweet broadcast to the world on one of the internet's largest social media platforms is somehow immune from criticism. This feels like a discussion I'd have with my children: "remember kids, anything you put on the public timeline is there forever and you must assume that anyone could see it". Own your tweets, people.

Now, that said, there are certainly degrees to which I'll berate someone publicly be that in an embedded tweet in a blog post or directly on a social media platform. I avoid being derogatory or making personal attacks and I focus instead on the topic being discussed. I'll debate that vehemently at times, but when someone chooses to have that debate on social media (and they do make that choice, it's not forced upon them), then it's now a discussion in the public domain and they'll need to bear the consequences of that, which might include my highlighting the tweet and others chiming in.

One last observation on this practice is what happens when someone deletes a tweet that's been embedded somewhere, as is the case with some of the ones I highlighted in this blog post. Just like I also teach the kids, deleting something from the internet doesn't make it permanently going away. Using Twitter's embed code, the tweet goes in as a script tags that paints an iframe into the DOM in a nicely formatted fashion complete with the profile photo of the tweeter, their name, date of the tweet, the number of likes and RTs and other associated metadata. If the tweet is deleted then all this disappears, however...

Twitter's embed code also includes a blockquote tag with the entire contents of the original tweet which means 2 important things:

- 1. Deleting a tweet doesn't remove the contents of the tweet from places it's been embedded
- 2. Deleting a tweet means in places where it's embedded, it's now clear that the tweet has been deleted

Which all kinda makes it look like someone regretted their tweet and later removed it, except that both the tweet and the regret are on full show to everyone. It's one of the reasons I avoid ever deleting a tweet I regret (and there's been a few, namely because I worded things poorly), rather I'd add to the tweet thread and clarify the position. That upsets some people too but hey, it wouldn't be Twitter without people getting upset!

EXTENDED VALIDATION CERTIFICATES ARE DFAD

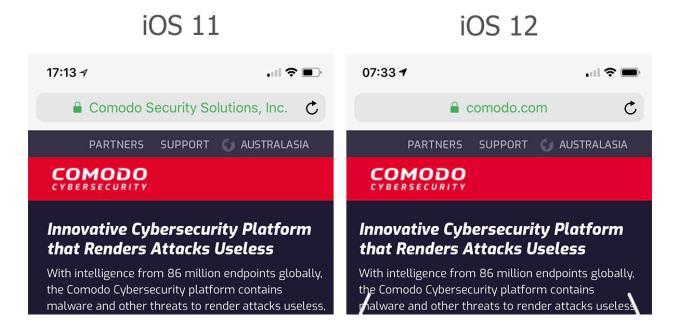
I have very little patience for bullshit. I suspect that behavioural trait is part me being Australian and part me being, well, just me. It means that when faced with the aforementioned bullshit, I find myself being quite intolerant and prone to sharing some pretty direct thoughts on the matter. This brings us to extended validation certificates, a topic I'd given a lot of thought to over many years and had emphatically concluded were useless. Possibly even worse than useless.

One of the things that really prompted me into action here was the shady marketing around EV. You'll see examples of that in the blog post and they were just a small sample of what I'd seen over the years. That it was coming from massively well-resourced companies like Comodo only served to motivate me more to write this post. But this one almost felt like it wrote itself because there was just so much material to support my position, hence my feeling the need to include the iconic Simpsons "he's already dead" meme. Because it was - and now it *really* is.

18 SEPTEMBER 2018

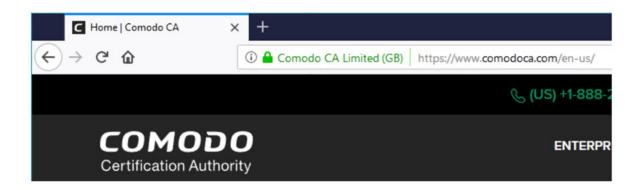
hat's it - I'm calling it - extended validation certificates are dead. Sure, you can still buy them (and there are companies out there that would just *love* to sell them to you!), but their usefulness has now descended from "barely there" to "as good as non-existent". This change has come via a combination of factors including increasing use of mobile devices, removal of the EV visual indicator by browser vendors and as of today, removal from Safari

on iOS (it'll also be gone in Mac OS Mojave when it lands next week):



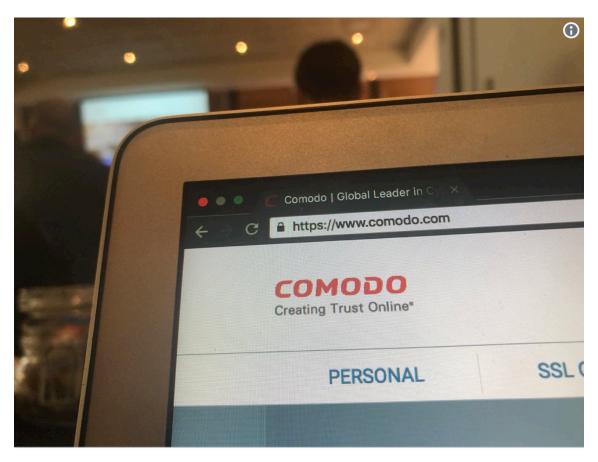
I chose Comodo's website to illustrate this change as I was reminded of the desperation involved in selling EV just last month when they sent around a marketing email with the title "How To Get The Green Address Bar On Your Website". The "alternate truth" of what EV does comes through very early on, starting with this image:

For over a decade, browsers have supported two different levels of security indicators for HTTPS pages: the standard "Secure" display and the expanded "green address bar" display like this:



The expanded green address bar is activated by an Extended Validation (EV) SSL Certificate. Here's how it works...

This is indeed what Firefox looks like today, but they entirely neglect to mention anywhere within the marketing email that this is an arbitrary visual indicator chosen at the discretion of the browser vendor. Obviously Apple have already killed it off, but even for many people on Chrome, the Comodo website actually looks very different:







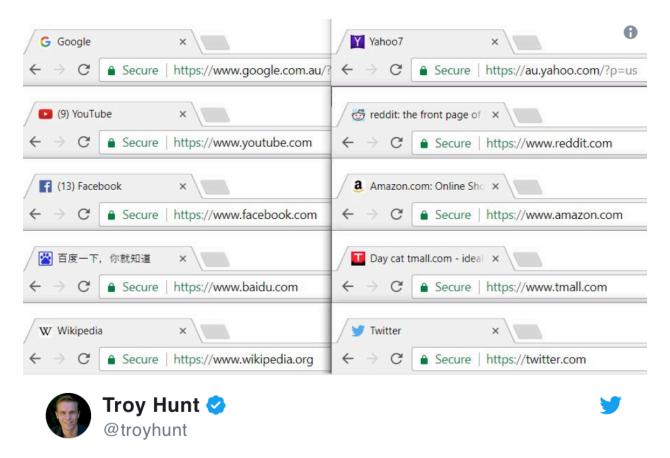
So it turns out that 3 different machines in my workshop today are part of the Chrome experiment to remove the EV indicator from the browser. The usefulness of EV is going, going...

The email goes on to talk about how EV fights deceptive websites and claims the following:

The verified company name display allows the user to quickly determine the legal entity behind the website, making phishing and deception harder.

In other words, seeing the company name results in higher levels of trust or if we invert that statement, *not* seeing the company name results in decreased trust, right? The problem is, people simply aren't conditioned to *expect* to see the

company name and there's very simple, effective demonstration of why this is the case:



Watching with amusement as @CertCouncil (backed by commercial CAs) pushes for pricey EV certs. Here are the world's 10 largest sites: no EV!

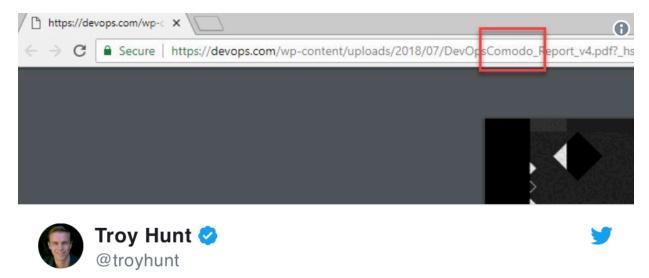
♥ 318 9:05 PM - Jul 13, 2017 · Gold Coast, Queensland

Comodo goes on with an attempt to establish the efficacy of EV by referring to "a recent study":

A recent survey by DevOps.com found that customers are 50% more likely to trust and purchase from a website with a green address bar.

They link through to a lengthy page on the Comodo store and whilst never explicitly saying it, use language that implies the study was somehow independent and unbiased: "Devops.com conducted a survey", and other such

phrases. <u>I shared a tweet thread about this back in July</u>, but this one tweet tells you all you need to know about the motives of the "survey":



Replying to @troyhunt

The first thing I wanted to know was who commissioned this and nowhere in the report does it tell you. However, if you look at the URL...

♥ 33 9:21 PM - Jul 26, 2018

I did honestly try to get clarity on the source of this work as well, first by tweeting the author of it then, after not receiving a reply, following up with him again and copying <u>@TechSpective</u> for whom he's the editor in chief along with <u>@devopsdotcom</u> (which follows me) who published the survey:





Still trying to get an answer on this, can @RealTonyBradley, @TechSpective or @devopsdotcom kindly clarify? The motives behind the study are important for reasons that should be apparent in the thread.

Does anyone else have a contact they can leverage to get clarity?

Troy Hunt @ @troyhunt
Replying to @BretCarmichael @Scott_Helme
Hi @RealTonyBradley, can you clarify who commissioned that

○ 6 8:42 PM - Aug 23, 2018

study?

0

Eventually, what was already abundantly clear was confirmed:



y

Replying to @troyhunt and 2 others

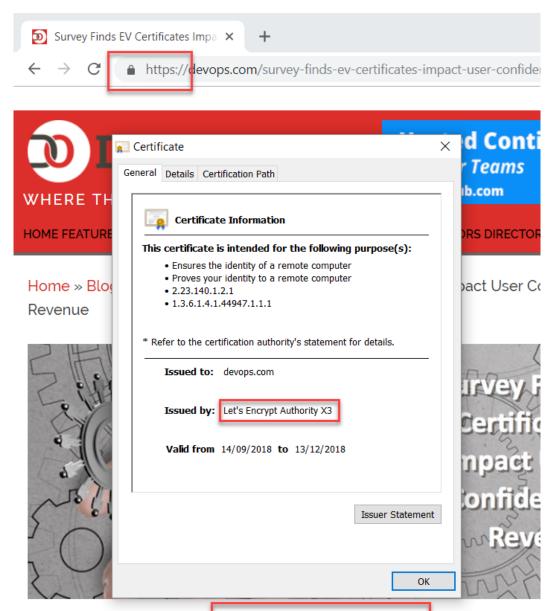
Hey. My apologies for taking so long to respond. I post a lot, but rarely look at mentions or replies on Twitter. The report was commissioned by Comodo CA.

◯ 1 11:53 AM - Aug 24, 2018

0

I wish this was made clear in the report itself because Comodo's vested interest is clearly going to introduce bias. It'd be like an oil company commissioning a report that concludes fossil fuels aren't harmful to the environment or a tobacco company stating smoking doesn't lead to adverse health outcomes. If you ever

had any doubt about whether DevOps.com actually believes in the "findings", take a look at how much confidence they themselves have in EV certificates and who they chose to go to when acquiring a cert:



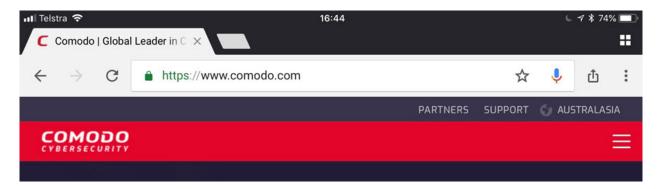
Survey Finds EV Certificates Impact Confidence, Revenue



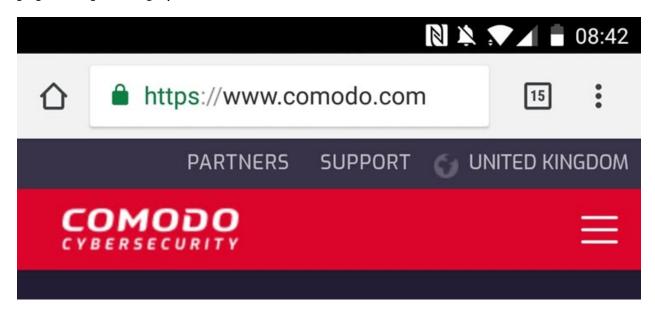
This resource is mentioned again throughout the Comodo email but we'll skip that for now. Moving on, they then state that you can "activate the green address bar" simply by purchasing an EV cert:

To activate the green address bar on your website, you just need to purchase and install an Extended Validation (EV) SSL certificate.

Unless you're using the world's most popular browser running on an iOS device:



Same again if you load the site up in Chrome on an Android, the world's most popular operating system:



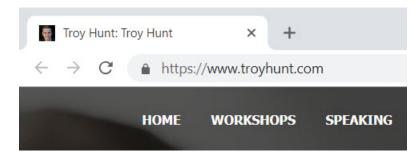
Even try going to Microsoft Edge on iOS and it's a now predictable result:



These are really, *really* important images as far as the value proposition of EV goes for two key reasons: Firstly, <u>we're approaching two thirds of all browsing being done on mobile</u> which means that those images above - *the ones that don't show EV* - are the predominant browsing experience any website owner should be considering. Secondly, as a result, this means that companies cannot tell their customers to expect EV because most of them will never see it. Despite this, Comodo suggests there's value in EV because of the "bigger security display":

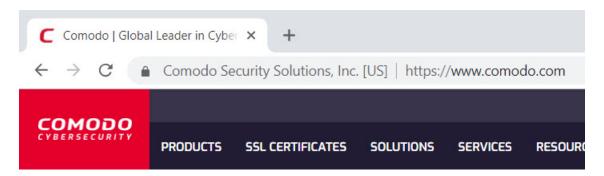
The larger security indicator makes it very clear to the user that the website is secure.

You know what makes people think the website is "secure"? When the website says "secure" just as it does next to the URL in the browser right now if you're reading this in Chrome on the desktop! Paradoxically, you only get the "secure" indicator when *not* using an EV cert and one could quite reasonably argue that this actually creates a greater sense of confidence by literally using the word "secure". And in case you're reading this and thinking "hang on, Chrome doesn't do that anymore", you're completely right:

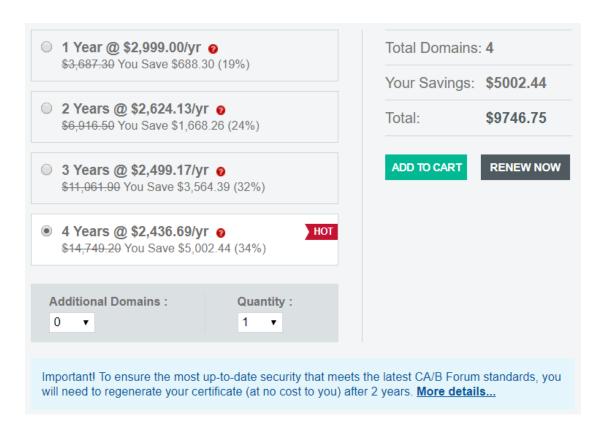


I wrote the first part of that paragraph before Chrome 69 hit on September 4 and

removed both the "Secure" text and the green indicators. That's not just a DV change either, sites with EV now also look rather different:



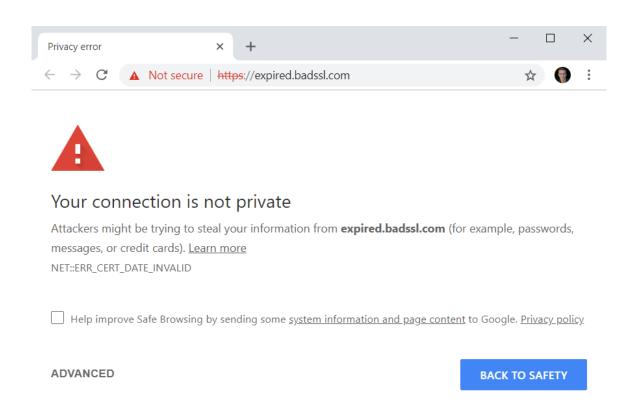
The point I'm trying to highlight here is both the fact that visual indicators are entirely at the discretion of the client and that they change over time. As such, the title "How To Get The Green Address Bar On Your Website" is now even more incorrect than it was when it was written! In fact, the only piece of the email that even came close to accurately representing EV was the admission that you can't get an EV wildcard cert. But wait! There's a solution and it's easily available just by spending more money, it's called a multi-domain certificate and the default option when looking at Comodo's Enterprise SSL Pro with EV Multi-Domain product will actually save you \$5,002.44*:



* Note: You must spend \$9,746.75 before the saving is realised

To be clear, this isn't a 4-year certificate either; as the text at the bottom of the image points out, the CA/B Forum guidelines limit certificate validity to 2 years and after that you need to manually go back through the entire verification and issuance process again. But hey, let's not allow that to get in the way of selling 4 year's worth of certs!

And what if you *don't* renew the cert then? Well, you get a great big pile of <u>this</u>:



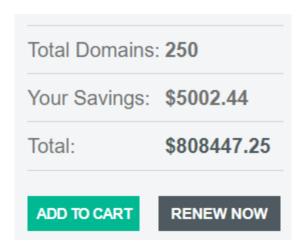
Now, you may be thinking "well that's kinda obvious and the same holds true whether it's EV or DV", but it's more nuanced than that. Firstly, neglecting to renew a cert happens with alarming regularity and it happens to the big guys too. For example, Microsoft failed to renew secure.microsoft.co.uk back in 2001. Too long ago? They also failed to renew an Azure one in 2013 and just to be clear about it certainly not being a Microsoft thing, HSBC forgot one in 2008, Instagram forgot one in 2015 and LinkedIn forgot one last year. There are many, many more examples and they all adhere to the same underlying truth; if something is important and repetitive, automate it!

Which brings me to the second point: certificate renewal should be automated and that's something that you simply can't do once identity verification is required. DV is easy and indeed automation is a cornerstone of Let's Encrypt which is a *really* important attribute of it. I recently spent some time with the development team in a major European bank and they were seriously considering ditching EV for precisely this reason. Actually, it was more than that reason alone, it was also the risk presented if they needed to quickly get

themselves a new cert (i.e. due to key compromise) as the hurdles you have jump over are so much higher for EV than they are DV. Plus, long-lived certs actually create other risks due to the fact that <u>revocation is broken</u> so iterating quickly (for example, Let's Encrypt certs last for 3 months) is a virtue. Certs lasting for 2 years *is not* a virtue, unless you're coming from the perspective of being able to cash in on them...

(Paradoxically, the LinkedIn story I linked to above is on TheSSLStore.com which is a certificate reseller. You can probably see where this is going, but rather than suggesting that automation is a key part of the solution to cert renewal, they instead suggest solutions "that scale to Enterprise level" from CAs such as Comodo who, of course, are pushing EV. No mention of Let's Encrypt, but then this is also the company that's been vocally critical of them for issuing certs to phishing sites (that do correctly validate domain ownership) whilst neglecting to mention that Comodo was issuing just as many at that time!)

A lack of wildcard support is one of the big *technical* reasons EV is avoided (the other reasons are mostly just common-sense ones), and loading up subject alternate names is a barely sufficient alternative. For example, we use a wildcard cert for <u>Report URI</u> so that you can send reports to https://[my company name].report-uri.com and we've got hundreds of those. Comodo will happily support that scale too:



Other than the fact that Scott Helme and I aren't really in a position to shell out

\$808k, this is also a far cry from what a genuine wildcard cert does as you need to specify all host names at the time of issuance as opposed to being able to dynamically serve them up.

The final point of note on the marketing email is the promise of a warranty:



Comodo's Most Popular EV SSL

Our most popular EV SSL certificate is the PositiveSSL EV, which is a very affordable EV cert with a generous \$1,000,000 warranty. Read more...

That actually links straight back to the page with the super pricey multi-domain EV certs and doesn't even attempt explain what the warranty is, which is a bit odd. But it's also consistent because <u>nobody actually knows what the warranty is and if anyone has ever claimed it</u>. Seriously - that's not intended to be a flippant statement, Scott and I genuinely tried to get to the bottom of that earlier this year and we simply couldn't get straight answers. When we did manage to engage in dialogue, I was accused of being in "nerdville":



Andreas Mallek @amallek · Feb 24, 2018



Replying to @andygambles and 2 others

Andy, those guys don't want to accept that they have a different point of view - they are way too nerdy to accept that normal people have different needs than people in nerdville. I am done tweeting in nerdville and back focusing to my customer's real world problems. Cu



Troy Hunt contract contr

I asked a very reasonable question Andreas and it's an important one because certs are being sold with a warranty and I'm trying to understand what that means. Real customers want to know this - what does it cover and are there documented examples of it being used? Do you know?

This was admittedly a very surprising response from someone that holds a position as the CEO at <u>CertCentre</u> because one would imagine that he, of all people, would want to espouse the virtues of cert warranties (assuming there actually are any, of course). If you're paying a company like CertCentre money for a product with a stated set of features, being a "nerd" by asking how those features work seems perfectly reasonable and not something that should result in ridicule from the bloke running the place. Unfortunately, rather than answering the question, Andreas decided it was easier to take the tried and tested <u>ostrich approach</u>:



Andreas Mallek

@amallek

@amallek blocked you

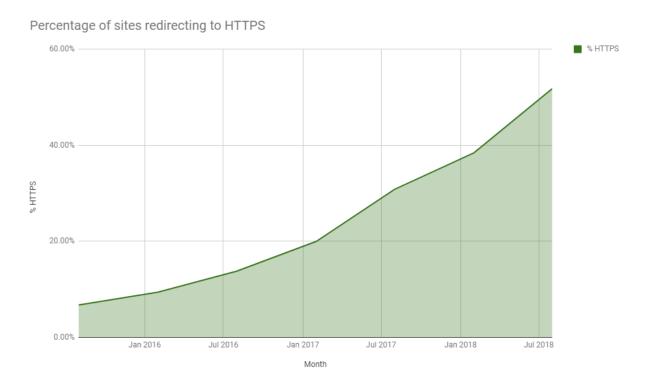
You are blocked from following @amallek and viewing @amallek's Tweets.

The thing I have a real issue with here is that there's a financial incentive to promote the warranty (you certainly don't get a warranty with a Let's Encrypt certificate), but no willingness to explain what you get for your money. CertCentre actively lists warranties as a "Top Security Feature" too:

Encryption Strength	Up to 256-bit SHA-1/SHA-256	Up to 256-bit SHA-1/SHA-256	Up to 256-bit SHA-1/SHA-256
Full Organization Authentication	✓	~	✓
Extended Validation	×	×	✓
Green Bar	×	×	✓
"Warrenty	\$1,500,000.00	\$1,500,000.00	\$2,000,000.00

But hey, if you can't even *spell* warranty, what are the chances of actually understanding what it does?!

Driving the nail even further into the EV coffin is <u>Scott's 6-monthly Alexa Top 1M report from last month</u>. In here he shared a very encouraging stat which is the growth in sites redirecting from HTTP to HTTPS:



It's now 52% which is enormously positive for the web in general. But it was this comment about EV which piqued my curiosity:

Despite seeing strong growth in HTTPS across the top 1 million sites, EV certificates have not seen much of that growth at all.

Let's put it in raw numbers: in Feb there were 366,005 sites redirecting from HTTP to HTTPS and 19,802 of them used EV certs so call it 5.41% of all HTTPS sites using EV. Fast forward to August and there were 489,293 sites redirecting to HTTPS with 25,158 serving up EV certs which equates to 5.14%. In other words, the EV market share declined by about 5%. As a proportion of all sites using certificates, EV is far from growing, it's actually going *backwards*.

(Incidentally, in case you're looking at the 489k figure above and thinking "that's actually less than half of 1M", Scott's scan failed on about 47k websites so they're excluded from the stats.)

As it turns out, many sites are actually *removing* EV certs. Last month <u>Scott</u> <u>detailed a number of major sites that used to have EV</u> and they spanned everything from Shutterstock to Target to UPS to Visa to the UK police. Around the same time, I noticed that even Twitter had killed their EV cert:

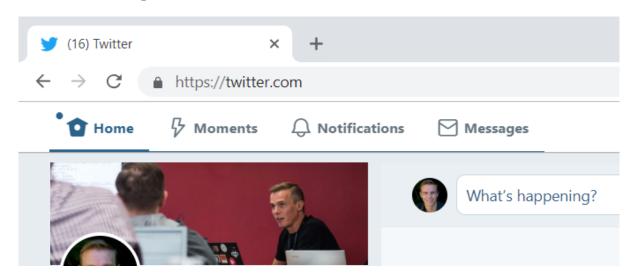


Hey, anyone else notice that Twitter recently ditched their EV certs? I'd love to know why (I mean other than the fact they're completely useless). @Scott_Helme?

Looks like the move kicked off a couple of months ago: crt.sh/?q=twitter.com

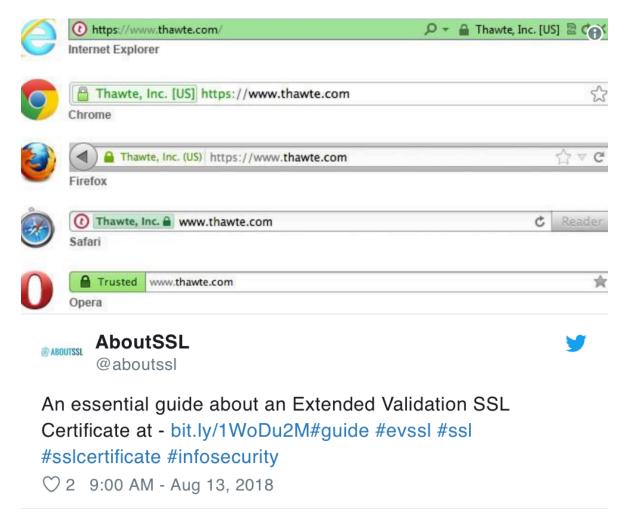
Twitter has been a bit of an odd duck for a while as far as EV goes; back in the earlier tweet showing the world's largest websites don't have EV, there were a bunch of replies from people saying it *does* have EV. We later discovered that depending on where you are in the world, you may or may not see EV on

Twitter. For example:



Certainly, as of today, EV is not being served up when I connect from Australia so for whatever reason, Twitter don't see it as important enough to show consistently and will switch in and out of EV as you move across the globe. That also says something significant about the effectiveness of EV: if they're willing to constantly add and remove it depending on where you are, do you think people are behaving differently and no longer trusting the site when they *don't* see EV? No, of course not, but that's the foundation that the mechanics of EV is built on!

I don't just want to focus on Comodo and CertCentre though because disinformation campaigns go well beyond those 2, for example:



Moving past the choice of historic browsers used in the illustration (just how old is that image?!), the piece that tweet links to makes the following claim:

Web security experts recommend adopting EV SSL Certificate for platforms such as E-commerce, Banking, Social Media, Health Care, Governmental and Insurance platforms.

Now I'm not sure who they're referring to in those first few words, but I do know that with the exception of banking, that statement simply doesn't hold water for the remaining industry sectors. It only takes a few minutes to demonstrate how fundamentally wrong this is so let's do it now:

Here's <u>the world's top shopping sites</u>, click through to see if any of them are on EV:

- 1.Amazon
- 2.Netflix
- 3.<u>eBay</u>

You might argue that Alexa has miscategorised Netflix as "shopping" so just for good measure, try the next largest which is <u>walmart.com</u> and, well, it's the same result. No EV. Anywhere.

Moving on and social media is the same deal:

- 1.Facebook
- 2.Twitter
- 3.LinkedIn

As discussed earlier, Twitter has a bit of an identity crisis in terms of whether it's in or out on the EV front so give the 4th largest a go if in doubt which is Pinterest.

Onto the world's most popular health sites and it's more of the same:

- 1. National Institute of Health
- 2.WebMD
- 3. Mayo Clinic

No EV. Nada. Zip. Not a single one.

I couldn't find one clear listing of global government websites so I pulled the data from <u>Scott's nightly Alexa Top 1M crawl</u> and grabbed the biggest .gov ones. The NIH was the largest but we've already covered that so let's take the next 3:

- 1.<u>Unique Identification Authority of India</u> (which has <u>other fundamentally</u> basic HTTPS problems)
- 2. Indian Income Tax Department

3.GOV.UK

By now you'll already realise the chances of EV being anywhere aren't real good. You're right - not a single EV cert to be seen.

Last up is <u>the top insurance sites</u>:

- 1.United Services Automobile Association
- 2. Kaiser Permanente
- 3.Geico

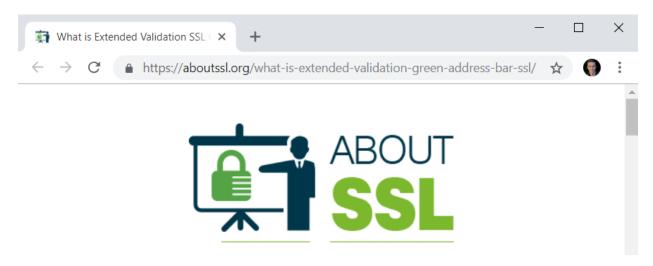
We got one! The USAA actually *does* have an EV cert! The other two don't but hey, at least that's something, right?

If "web security experts" are recommending EV for sites of these classes then clearly those responsible for actually making the decisions aren't listening. Except that nobody who's actually thought through the logic of EV properly is actually making these recommendations anyway so perhaps there's just a bit of poetic licence there in the copy.

Another set of unsubstantiated claims made by About SSL is that EV "increases transaction conversion rates", "lowers shopping cart abandonment" and "protects from phishing attacks". You can understand *why* they're making these claims and there's a pretty clear call to action immediately under the list of conveniently bold green selling points of EV:

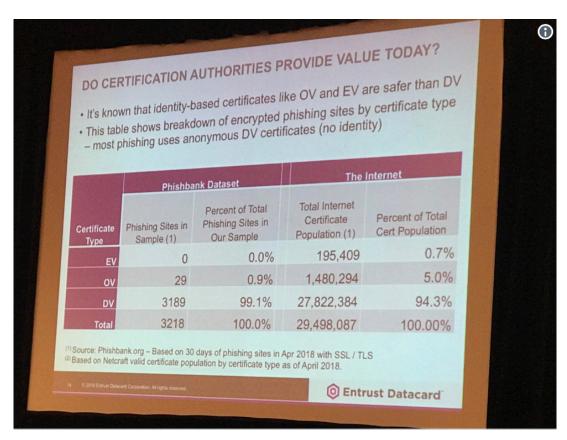


So we're back to there being a clear bias again. But hey, they're just out there trying to run a business so I get the motives. One would also assume that in running this business where you can purchase items online they'd like to increase their transaction conversion rates and lower shopping cart abandonment, right? Well there's a funny thing about that:



Even the company selling EV is smart enough to know it's not worth actually paying money for! Plus, of course, the whole "green address bar" thing is now completely defunct courtesy of the world's most popular browser killing it in version 69.

But then there's the phishing situation and indeed this is often touted as being a strength of EV in that it somehow reduces it. In fact, this (much maligned) slide by Entrust from earlier this year makes precisely that point:





Adam Powers @apowers313



Replying to @iangcarroll

I think this was the basis of his argument for EV stopping phishing.

There's a whole pile of things wrong here and the best way to understand precisely what is to read through this thread from Ryan Sleevi who analysed the paper the claims were based on:



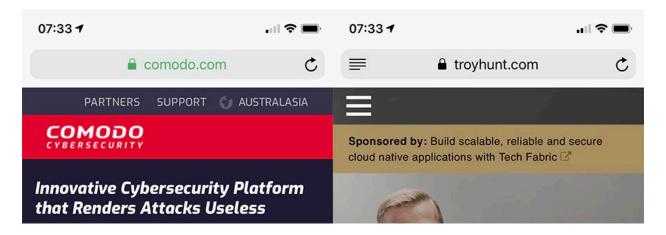
Ryan is a super smart crypto guy working on Chromium and has a very articulate way of tearing bullshit arguments to shreds. Towards the end of the thread he summarises the problem:



Replying to @sleevi_

10/ Look, it's a bad paper, but what's worse, is they try to pass it off as a "data-driven" research, using flawed methodology and cherry-picking to support a business model that relies on users having all the liability and responsibility for detecting UI changes.

And we're back to EV only being effective if people behave differently due to a UI change they don't know to look for and increasingly, doesn't even exist anymore. Either that or it's changed in nuanced ways people don't expect to look for; remember the first image in the blog post showing Comodo in Safari no longer displaying the registered business name in their EV cert? Take a look at it next to this blog, also loaded in Safari on iOS 12:



See the difference? The URL of the EV site and the padlock next to it are now in green whereas the DV site is in black. So now if you want to set an EV expectation you have to tell customers to look for the *green* URL and padlock... unless they're on Chrome which has now removed all the green bits! You can see how ridiculous this whole premise of telling normal everyday folks what nuances to look for in the browser is, especially with the rate at which they're changing.

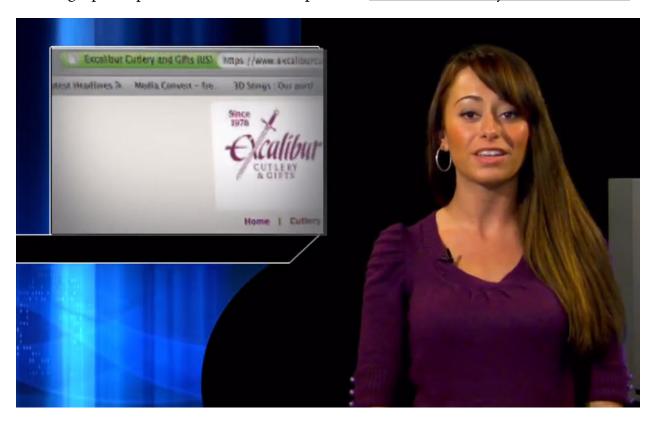
Back on the About SSL site, there's an embedded video which espouses the virtues of EV along the same sorts of lines we've seen already. It's about 6 minutes long if you've got the patience to view it:



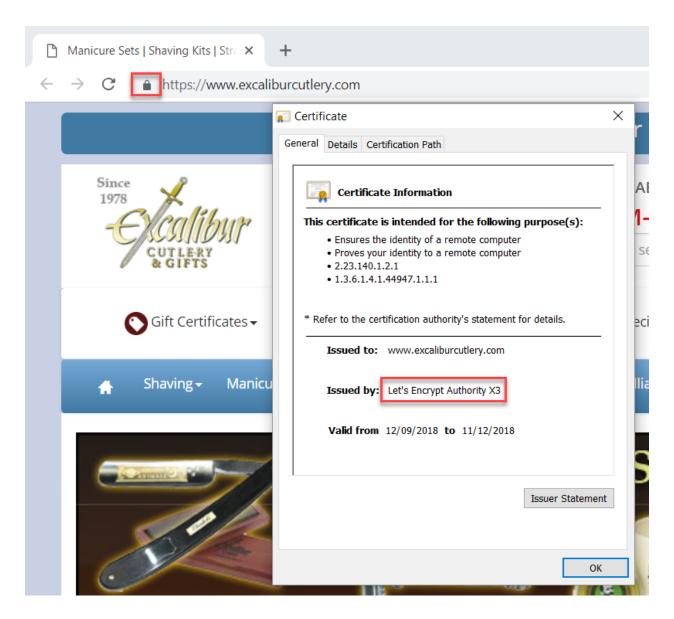
Or we can just skip to the good bits, such as when the presenter (and <u>Comodo Product Marketing Manager</u>) talks about the criticality of EV during a financial transaction:

Right at the moment of truth, when they're weighing whether or not to go forward with a transaction, this striking visual indicator (the green EV bar) accompanied by information certifying their business name, location and certification authority that validated it is presented providing needed reassurance to continue

Backing up her position is a screen cap of the Excalibur Cutlery & Gifts website:



You can probably sense where this is going by now... and you're right:



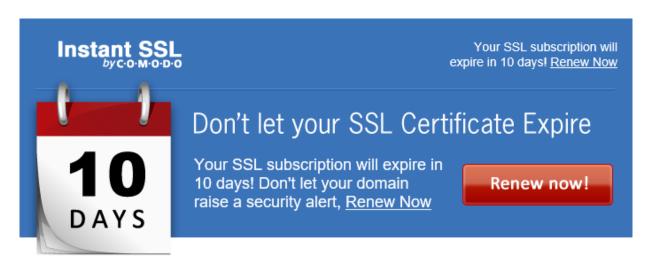
No EV. No commercial DV either but instead a perfectly good *free* Let's Encrypt cert. It's like the video was a remnant of a bygone era and as it progressed and showed websites running in IE8 on Windows XP I couldn't help but feel the information was somewhat... dated. Which turned out to be a fair assumption:

www.comodo.com

© Copyright 2009 Comodo Group, Inc.

Now I wouldn't normally hold a video of almost a decade ago against today's standards were it not for the fact that the views expressed there are consistent with those expressed today. Plus, of course, the video was linked to from a tweet only last month under the guise of "An essential guide about an Extended Validation SSL Certificate" so it's fair game in this case.

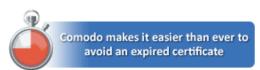
Comodo using sites to promote EV that *don't* use EV seems to be a bit of a pattern. Just this month, someone forwarded me on a domain renewal email they got from Comodo that looks like this:



Renew in Seconds

Renewing your SSL certificate has never been easier. Take a second now to make sure your trusted website stays that way.

Renew Now!

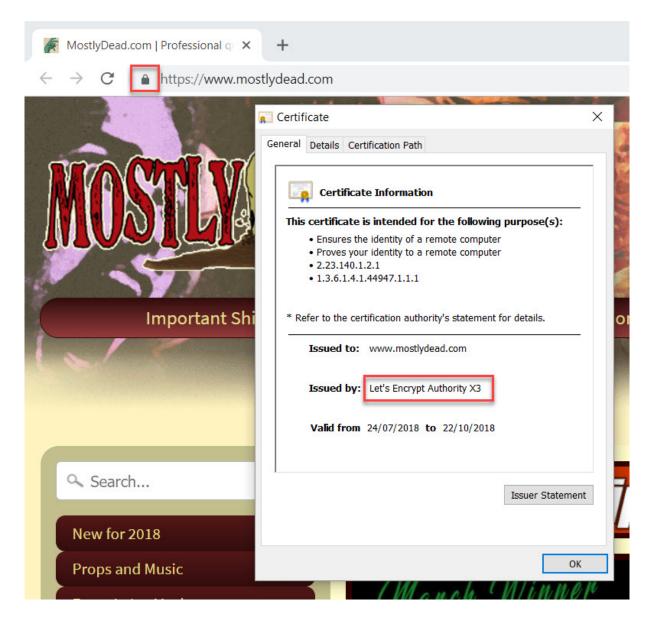


Boost Online Sales with EV SSL

Boost your sales with the trusted green address bar. "This product really creates consumer confidence. My sales increased 20%!" - Ken Kriz, Mostlydead.com. Learn more



Naturally, he was curious about <u>Mostlydead.com</u> and headed over to take a look at how well that "20% increase in sales was going". You know, because of how much EV "creates consumer confidence". Apparently, not so much anymore:

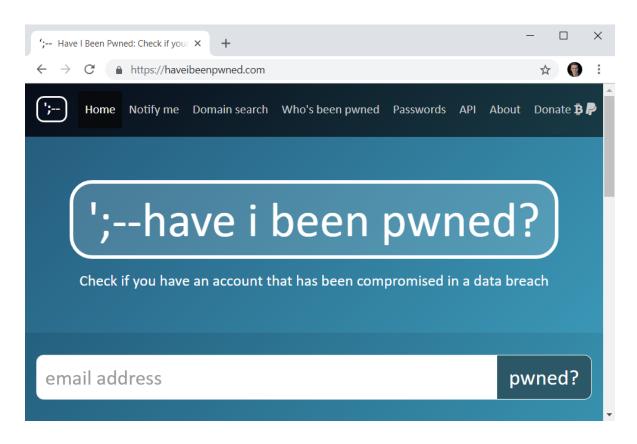


The more you delve into it, the more you can't help but conclude that EV is... mostly dead (we're beginning to see a pattern here). The thing is, this isn't just some random site that went from EV to DV, it's one that *Comodo specifically chose to show the value of EV!* This is meant to be a poster child site for the value proposition of extended validation and it's one Comodo still promotes to this very day. Yet, here we are, with Ken Kriz obviously having a change of heart on the efficacy of EV (or possibly never having really been endorsed in it in the first place).

Right about now, the whole EV thing may be starting to feel a bit like this:



But we're not done yet, there's more and that brings me to another site which used to have EV and has now gone back to DV. It's this site:



I changed that cert just over one day ago and so far, nobody has even mentioned it. Nobody. Not a single person and I've got an audience that's far more aware of this sort of thing than your average person. There's certainly been no shortage of people that *could* have noticed it over that period too:



Nearly 2 years ago now, I wrote about <u>my journey to an EV cert</u>. Like many of the posts I write, this one was as much for my own education as it was for yours; I wanted to go through the EV process myself (it had always been done by other teams in my previous roles), and frankly, I wanted to see if it actually provided any value. I honestly didn't know at the time and I summarised the post as follows:

This whole EV cert thing is hard to measure in terms of value; I have no idea how many more people will put their email address into HIBP or how much more media or good will or donations it will get. No idea at all.

A couple of years on, I'm pretty convinced of the value: there isn't any. Now that's not to say there was a *downside* to having the cert in place as I became increasingly disillusioned with the whole premise of EV, but rather there's also no upside. As the renewal date approached (it was 14 December), I made the call to proactively kill the cert and roll over to a free one issued by Cloudflare. There was absolutely no reason at all to pay the renewal fee (I'd previously paid \$472 for a 2 year cert) and there was also no reason to wait to roll over to DV short of loss aversion which makes about as much sense as, well, EV certs.

I've often pondered the rationale of paying for EV certs and indeed paying for certs at all in an era of freely available ones. I spend a lot of time in companies around the world talking about HTTPS and when I probe on the decision-making process for certs, the phrase "nobody ever got fired for buying IBM" regularly comes up. I wanted to find a good reference to explain the intention of this phrase and I found an excellent one on Wikipedia's definition of FUD:

By spreading questionable information about the drawbacks of less well known products, an established company can discourage decision-makers from choosing those products over its own, regardless of the relative technical merits. This is a recognized phenomenon, epitomized by the traditional axiom of purchasing agents that "nobody ever got fired for buying IBM equipment". The aim is to have IT departments buy software they know to be technically inferior because upper management is more

likely to recognize the brand.

In other words, people are making uninformed decisions on what they think is a "safe bet" due to the marketing FUD. I suspect it's a similar mentality to companies placing third party security seals on their websites; they lack the sophistication to realise they can actually increase risk but hey, they were marketed well!

So that's it - EV is now gone from HIBP and nobody will miss it which would be entirely consistent with the experiences of others who've dropped it:



Replying to @troyhunt and 3 others

We dropped our EV this month, improved TLS handshake speed, and no one single feedback came in saying they missed it

◯ 3 8:58 PM - Aug 23, 2018





Replying to @spazef0rze and 2 others

We have replaced EV cert by @letsencrypt on our payment portal for this reasons:

- automatic renewal (no long and complicated manual process & reduces risk of expiration)
- price
- people don't care about cert type
- shorter expiration=quicker restore from potential compromise





0



Anthony Green @nthonygreen

Replying to @troyhunt @Scott_Helme

We realised that people trusted a nice green Secure more than our unfamiliar company name. Cost saving is a bonus

○ 3 5:45 AM - Aug 24, 2018



I disagree with that it's about cost. Target and the like don't give a crap about spending \$1k on a cert. I think it's about awareness. I know on my org's site, EV seemed like a good idea 18 mo ago. But will I renew it? Nope! Because I've learned just how pointless they are.

-- Jim Michael (@jimmichael) August 24, 2018





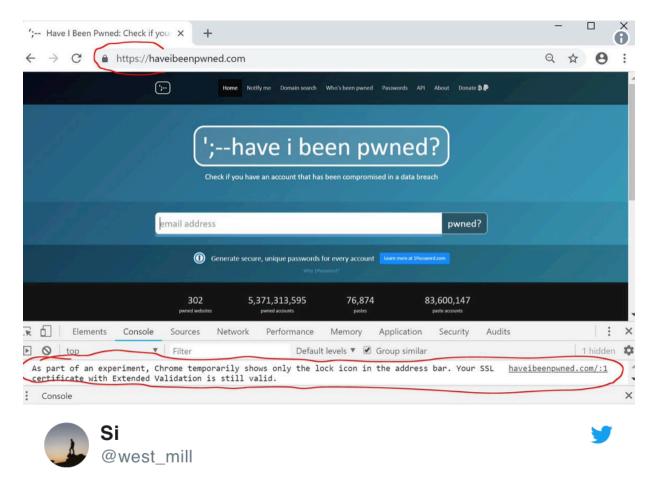
Replying to @troyhunt @Scott_Helme

Can't say who/where - biggest factors were 1. Needed a wildcard for improved flexibility. 2. Costs were no longer justifiable, especially considering multiple subdomains. 3. Lack of consumer awareness meant hardly anyone noticed the change.

◯ 1 9:25 PM - Aug 24, 2018



This turned out to be a long blog post because every time I sat down to write, more and more evidence on the absolute pointlessness of EV presented itself. I started jotting notes down well before some of the events listed above, not least of which was Chrome 69 and the removal of the green address bar which killed one of the big EV marketing headlines. It's hard to conclude anything other than EV has gradually suffered death by a thousand cuts; it was something that could be sold at a point in time in the past when the landscape was very different but today, it's just become a pointless relic of a bygone era. Browser vendors know this and are acting accordingly and it's only a matter of time before the final nail is in the coffin:



I see Chrome Canary v70 is trialing the removal of EV-SSL company names, wonder if that will make it into the final release? @troyhunt

That tweet was obviously from *before* I removed EV from HIBP and it's a glimpse into the future. When Chrome does finally remove the EV visual indicator from the browser (just as they've already done on mobile devices and as Apple has done across the Safari line), that'll well and truly be the end of EV. Perhaps then, the FUD will finally end.

I'll leave you with one final piece that explains the absolute futility of EV and it's a talk I did in London earlier this year. It's embedded at the point where I begin talking about EV and it's the audience interaction here that really makes it. Have a look at how a room full of smart technical people responds when I ask about

what visual indicators they expect to see on popular websites. Enjoy!



Comments

Not that I'm defending EV, but you can re-key a Digicert EV without going through any visible re-authentication.

Scott Helme: The validation can only be used for 13 months, after that you'd have to re-validate to change anything in the cert.

Troy: Well, that's an interesting one and I didn't want to delve into it in the main post but I will here since it's been raised:

Yes, you can re-key without going through the verification process again (although you'll need to go through the process again at renewal time), and I

actually went to do this just yesterday. I wanted to ensure that I could quickly restore the EV cert on HIBP if I needed to roll back for any reason, but I couldn't locate the key. I went to DigiCert to login and request it to be re-keyed but my creds didn't work. I went through the password reset process and never received the email. I then thought I'd go to CertSimple instead as they originally issued the cert so I logged on, lodged the request and waited for the email the UI said was on the way. It never arrived so I went through the process again and still never saw the email in either inbox or junk. In the end, I managed to locate the key and didn't need to pursue the issue any further. It was unequivocally my fault that I misplaced the key, but the whole process reminded me of the additional risks you take on with an EV cert. This was precisely the position the European bank I mentioned was taking; EV creates other associated risks that aren't there in the DV world.

I apologize if I'm being dense here, but it's not clear what the additional risks are with EV as compared to DV, from the above paragraph. There was risk introduced going through a reseller; there was risk introduced by using password based credentials instead of something like SAML SSO; there was risk introduced by using manual processes to issue and manage the EV certificate. I'm not clear on what components of risk were introduced with the use of EV which wouldn't be present with a DV certificate managed in the same way.

How far off am with the takeaway which I do have, namely that the approachability and availability of automation through systems like ACME and companies like Cloudflare abstracts away those risks to a significant extent (which is *huge*). Further, due to the general prevalance of DV certificates within those systems that allow for abstracting these risks, a correlation can be made that DV itself reduces the risk, but I think there's a minor disconnect between the conclusion and its underlying reasoning. i.e. both EV and DV certificate management and provisioning can be automated to approximately the same extent.

The risk is that, if your certificate expires without you having planned a new one then a DV

cert is nearly instantaneous whereas an EV certificate is a lengthy process so you have to make some plan for what happens while you are renewing your EV certificate.

That's definitely a valid point; it's not *nearly* as "out-of-the-box" to robustly automate EV renewal/replacement processes.

Troy: With a DV cert from Let's Encrypt, there wouldn't have been any requirement to re-authenticate to a service, I just would have issued an entirely new cert after proving I controlled the domain. I could have gone from zero to DV in moments instead of all the mucking around required to regain access to the account on the CA.

Epilogue

Time was not good to EV. 11 months after this blog post, I followed up with "Extended Validation Certificates are (Really, Really) Dead" in the wake of both Chrome and Firefox announcing they were dropping the visual indicator that was really the only thing that set EV apart from DV in any meaningful way. I do admit to feeling a bit smug about my prediction from the title of the earlier blog post coming true.

At a more macro level, one of the topics I kept finding myself drifting back to as my blogging and speaking career progressed was the impact of security controls on humans. Passwords, 2FA and obviously extended validation certificate indicators just to name a few; how do humans engage with these? I mean they're technical controls, but they rely on humans in order for them to work, so as well as looking at the mechanics of those controls we have to look at how humans work. I still find that idea fascinating and it extends well beyond security too. I later got very involved in IoT and found myself continually drifting back to that same question: how well does this work when you put it

into the hands of normal everyday humans? Very often and as with EV, the answer was simple - it doesn't.

10 PERSONAL FINANCE LESSONS FOR TECHNOLOGY PROFESSIONALS

I'd wanted to write this blog post for a long time - years, in fact - but it was one of those ones that took a lot of thinking about how to present the information before I felt ready to post it. I wanted to be transparent, but I didn't want to share anything too finically personal. I wanted to inspire others, but I didn't want to seem aloof. But like pretty much every other blog post before it, this was one that I just wanted to get off my chest and written up and finally, I'd dropped enough bits and pieces into it to finish it up and push it out to the world.

I carved out my niche online by writing about technical things whether it be programming or data breaches or just other cool tech stuff I enjoyed playing with. This blog post is none of those things, rather it's much more personal and much harder to clearly define than something like code. It's such an essential topic and one that touches every single person in one way or another, yet it's somehow so poorly understood. I wanted to include it here because to this day, I still have people telling me how it made a positive impact on them which is just wonderful to hear. It was also posted right around the time there were some major personal changes occurring in my life, which makes it all the more interesting to look back on now.

31 DECEMBER 2018

Patience. Frugality. Sacrifice. When you boil it down, what do those three things have in common? Those are choices.

Money is not peace of mind. Money's not happiness.

Money is, at its essence, that measure of a man's choices.

his is part of the opening monologue of the Ozark series and when I first heard it, I immediately stopped the show and dropped it into this blog post. It's a post that has been many years coming, one I started drafting about 5 years ago. One I kept dropping little bits and pieces into as the years went by but never finished because the time just wasn't right. It was only after reflecting on the responses to the following tweet that I decided to sit down and finally wrap up this post:





New family car! I'm kinda a bit excited about this one!



12:16 AM · Dec 24, 2018

This is a measure of my choices. Of my wife's choices. Of a couple of decades of choices. The car itself is only one small part of that measure, but it was the enthusiasm that tweet was met with by many who expressed a desire to do the same one day that prompted me to finish the post. It's also the negativity expressed by a small few that I should choose to spend our money in this way that prompted me to finish it; those that feel success itself or its manifestation into physical goods is somehow taboo. The latter group won't get anything useful from this post, but it was never meant for them. It was always meant for those who wanted the measure of their own choices to look more like the one above.

So here it is - 10 Personal Financial Lessons for Technology Professionals.

Intro: This Industry Rocks!

I want to start here because this post is very specifically targeted at people working in the same industry as I do. There'll be many things which I hope are useful to those outside of that, but frankly, those of us in tech have a massive advantage when it comes to our ability to be financially successful. I don't just mean at the crazy rich end of the scale (4 of the world's top 10 richest people did it in tech - Bezos, Gates, Zuckerberg and Ellison), but at all levels of our profession. In fact, those guys are all pretty good examples of the ability to build amazing things from the ground up and I'm sure that many of you reading this have sat down and started building something with the same enthusiasm as, say, Zuckerberg did with Facebook in 2004. Of course, success at that level is exceptionally rare, but my point is that in this industry more than any other I can think of, we can create amazing things from very humble beginnings.

But of more relevance to most of us is the opportunities this industry affords the masses. It's one you can get involved in at almost any age (<u>I started both my kids coding at 6 years old</u>), it provides endless opportunities to learn for very little or

even free (the vast majority of my own programming education has come via free online resources) and it transcends borders and socioeconomic barriers like few others (think of the opportunities it grants people in emerging markets). It's also up there with the highest paying industries around. I think we all know that innately but it's worth putting into raw numbers; I pulled a report from July put together by Australia's largest employment marketplace (SEEK) which has some great stats. For example, the ICT industry (Information, Communication, Technology) was the 5th highest paying with an average salary of \$104,874 (dollars are Aussie, take off about 30% for USD). Number 1 is "Mining, Resources & Energy" which had a local boom here but is now rapidly declining (down 14% on the previous year). Take mining out of the picture and the top industry ("Consulting & Strategy"), pays only 5% more than tech. Look the other way down the list and the next highest industry is "Legal", a whole \$9k a year behind. Banking is below that. Medical even lower.

Then there's this:

Today, the Information & Communication Technology (ICT) industry dominates, with salaries from six roles within the industry featuring in the top 20.

The highest salary SEEK has on the books is for architects (the tech kind, not the construction industry kind) at \$138k. The third highest is tech industry management roles at \$132k. Of course, actual numbers will differ in other parts of the world and indeed across other reports, plus there are many roles in the industry that will pay much less than those (especially during our earlier years). The point is that the tech industry provides people with near unparalleled earning potential across one's career. And it gives them the ability to do so much younger in life than many others do and with much less formal education; I care much more about skills than degrees in tech people, but my doctor / lawyer / pilot better have a heap of formal qualifications from many years of study behind them!

This is a cornerstone of what I'm going to write in this post: technology

professionals have a much greater ability to earn more than most other industries and to do so at a young age. Being smart with that money early on gives them an opportunity to leverage it into even greater things again. Keep that in mind because I'll come back to it in lesson 2 but firstly, let's just be clear about why all this is important.

Lesson 1: Money Buys Choices

Let me be crystal clear about this in the very first lesson: money is not about owning a Ferrari and living in a mansion. It's not about expensive jewellery and designer clothes. No, money is about choices. It's about having choices such that you can decide to spend it on what's important to you. That may mean helping out family members, donating to local charities or retiring early so that you can spend more time with your partner and kids. And yes, if it's important to you, it may also mean spending it on luxury items and that's fine because that's your choice! It's a choice you get to make with money as opposed to one that is forced upon you without it.

Let me share some examples of what I mean from my own personal experiences and I hope they cover a broad enough spectrum to resonate with everyone in one way or another. Just over 2 years ago, my wife (Kylie) had spinal surgery. You can read her experiences in that post but in a nutshell, it wasn't much fun and it followed many months of pain due to disc degeneration. The choice that money gave us was to focus on her treatment and recovery without stressing about her needing to work. We said to each other many times "how on earth would we have dealt with this if she still had a full-time job?" and invariably the answer is always that we couldn't have: the job would have gone.

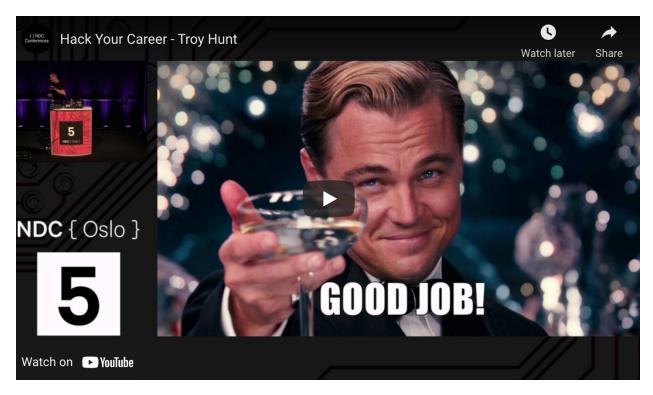
Kylie wasn't working when her back went because we chose not to. She left a very successful corporate role in late 2014 and very shortly after, <u>my own corporate job was made redundant</u>. We never really consciously decided that she

shouldn't go back to work, but a series of events including her being fed up with corporate life and us deciding to move interstate meant that she never did (although she's continued consulting on an ad hoc basis). Money gave us that choice. It was a choice that meant one or both of us is always there for the kids in the morning, always waiting to pick them up after school and always there for every tennis match, friend's birthday party or other random kid thing that seems to happen on a near daily basis. Being able to make those choices has enabled us to spend more time together as a family. It's quite literally bought us family time in many different ways, particularly in recent years.

Which leads me to the "but money can't buy happiness" position so many people have repeated over the years. Bull. Shit. Anyone who has ever said that simply doesn't know where to shop. Putting aside the intangible things money buys such as those already covered above, money spent on physical items can bring people a huge amount of pleasure. I'm not a fashion guy (pick almost any talk I've done and you'll see it's jeans and t-shirts all the way), but I totally understand how presenting well can bring a lot of joy to people. Obviously I am a car guy and vehicles such as the one at the beginning of this post and the Nissan GT-R I bought back in 2013 have brought me enormous pleasure. I smile every time I drive either and the latter in particular has resulted in so many immensely enjoyable interactions with people; kids taking pictures, adults wanting to chat and without exception, positive responses from everyone who sees it. Now mind you, some of the most fun times I've had have been in previous cars a fraction of the price so I'm by no means trying to imply a direct correlation between cost and happiness, the point I'm making is simply that tangible items that cost money can bring a huge amount of happiness, but only if you have the choice to obtain them.

I'm very conscious of the fact that for some people, signs of wealth lead to resentment. There was some of that in response to the Mercedes tweet earlier on and in Australia, we'd refer to that as <u>tall poppy syndrome</u>. (I'm still at a loss as to why anyone would take the time to explicitly tell you how displeased they are with your happiness; some people just lose their minds when they're behind

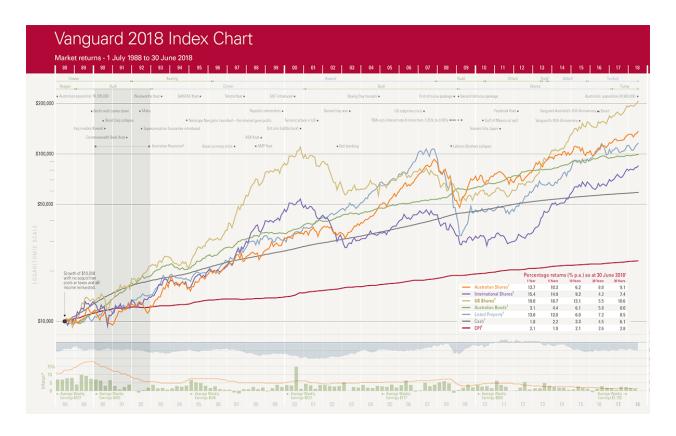
a keyboard.) I also touched on this when I first did my Hack Your Career talk in Norway last year where they refer to it as <u>Janteloven</u> (video embedded at the point where I describe it):



For the purposes of this first lesson, I don't care whether someone feels this way or not but regardless of your position, the one thing you should take away from this is that money enables you to choose what's important to you, whatever that may be. That's the mindset you need to take as you progress through this post.

Lesson 2: The Money You Earn Young is the Most Valuable Money You'll Ever Earn

Let's start with a graph and it's one you may have seen before, or at least some interpretation of the same sort of data:



This is <u>Vanguard's 2018 Index Chart</u> and you can either drill down into it and pour over the details or just take one simple truth away from a glance at it: investments grow over time. I know, revolutionary, right? Now to be fair, some investments tank and others skyrocket but what's more important than the minutiae is the overall market forces that enable money to multiply over time. We're looking at 30 years here and \$10k invested back in 1988 would be worth almost \$59k today invested at cash rates (6.1%), nearly \$85k if put into international shares or over \$206k if invested in US shares (and that includes the GFC period). There's also CPI at work which makes that \$10k worth less today than it was 3 decades ago, but that's tracked at 2.8% per annum which is a damn sight less than a balanced portfolio earns.

An often-heard saying illustrates the value of starting early and allowing time to amplify investments:

It's not timing the market, it's time in the market.

In reality, it's both and buying anything at a low-point is obviously going to net

you more dollars than buying at the peak. But the point of all this is that starting young enormously amplifies earning potential and to bring this back around to the tech industry again, those of us in this space have a much better chance than most to earn well at a young age. Let me put some personal context around this:

Almost 9 years ago I wrote a post on a real estate forum looking for feedback and inspiration. We'd bought a lot of property by that time and it became the foundation on which so much of what we've done since has been built. This is the first time I've mixed these two worlds - my background with real estate and my public blogging life as most readers will know it - but it's important context. Do read that post as it goes a long way to explaining why I'm writing this post and indeed, why I have the financial options I do today.

Kylie and I started investing while we were young. We began purchasing real estate in 2003 in our mid-20s and we poured every cent we could save into it. Some purchases were better than others, of course, but the constant theme across all of them was that we knew that good investments made young would pay off big time in the long term. It also created a forced savings plan for us; money in real estate is not "liquid" so you can't readily draw it out of a savings account on a whim and loans need to be paid on time each month or banks start getting cranky. (Incidentally, this is also a strength of home ownership as it's effectively a forced savings plan.) We maximised our borrowing potential, took advantage of every available tax concession and relentlessly pursued more property as soon as we had the savings to put down another deposit. We took risks, but they were calculated and made at a time where we had 2 incomes and no dependants. Everything gets harder when there's kids; more expenses, less time and often, less income if one partner decides to stay at home or work less.

By no means am I saying "go out and put all your money into property", it might be that you start putting a very small amount of money into a share portfolio or managed funds early on, the point is that time amplifies money (at the very least, everyone should understand how <u>compound interest</u> works). That was the single best financial decision we ever made and it happened well before my life

as people know it today; there was no Pluralsight, no workshops, no speaking events or Have I Been Pwned or blog sponsorship - nothing. Yet today, that property portfolio is a significant portion of our wealth because even though we weren't earning much money then by comparison, it amplified over and over again.

I want to touch on 2 more things on this because I know they'll come up if I don't mention them. Firstly, if you've passed the age that you might consider "young", the same logic of time amplifying dollars still applies. Obviously, you have less time and there are other considerations such as retirement funds (and associated tax implications), the point is that the earlier you begin on this journey, the better. And secondly, no, this wasn't done with financial support from parents. No deposits were handed out, no financial guarantees were made on our behalf, every single cent had to be earned, saved and then invested. But there was some help we got that moved everything along, and that was with financial literacy.

Lesson 3: Invest in Financial Literacy

I regret many things about my own education at school and university. I regret that I had to learn French in high school. I regret that I had to do chemistry as part of the computer science degree I started and never finished. But most of all, I regret that I was never taught financial literacy. I never learned the importance of the things I've already written in this blog post nor how the share market or property market work or even something as simple as the impact of compound interest on a credit card, something that's at crisis level for many people here in Australia at the moment. These things, to my mind, are essential life lessons and I do hope things have moved on a bit in the education system since then.

But we did have encouragement from our parents when it came to imparting financial advice. The two most notable things that come to mind were my father regularly repeating lesson 1 above (money gives you choices), and Kylie's father helping us understand how the property market works (he worked in the industry). But that was a tiny portion of the education with the vast bulk of it made up of reading books and magazines, going to seminars, hanging out on forums and frankly, also learning by making mistakes. We lost money on shares. We missed opportunities that would have yielded amazing results. We had property deals fall over. We got a lot of stuff wrong, but we got a lot more stuff right.

Part of developing financial literacy is that the more you learn about money, the more conflicting advice you'll get. Last week I tweeted about drafting up this post and I had a number of people contact me with their own tips. One person emailed me with many that aligned with mine, but he also said "only buy properties that you feel you could live in, they are homes as well as investments". I would never want to live in any of our properties we bought as investments. When you buy an investment - any investment - you should be ruthlessly focused on the numbers; what it's yielding, what the growth opportunities are, the tax advantages etc. When you buy a home to live in, you're buying with the heart because a home is a very emotional purchase. That's not to say you can't buy a home that's also a good investment, but you have different priorities and the perfect home for you to live in is almost certainly not the perfect asset for you to invest in. I don't want to live in any of our properties, but they're in high growth areas with good accessibility to public transport and low vacancy rates. Now, that doesn't make me right and him wrong, it's merely an illustration that there are many different views out there and the challenge for you is to understand the reasoning behind them and work out what actually makes sense for you. That knowledge is an investment you have to make.

Financial literacy is a fundamental skill which we all need but few of us genuinely invest in. There are a heap of resources available where you can learn for free and whilst there's frankly a lot of crap out there (the are way too many dodgy characters trying to sell investment opportunities!), it all contributes to

the melting pot of information you can absorb. I'm conscious that for most people, developing financial literacy probably seems like a difficult thing that requires a time commitment. And I agree. I found it hard and I found a huge amount of my time being spent on it, but I do believe that we, fellow geeks, have some advantages here.

Those of us in the tech industry are used to seeking out information online. Crikey, I still use Google every time I need to write text to a file in C#! We're also used to engaging with others online in order to learn, we've been doing it on Stack Overflow for years and we can do it on any number of investment forums, debt support communities or other resources designed to help educate in the same way as the tech ones we're so dependent on. I made 414 posts on the property forum I referenced earlier, more than all my questions and answers on Stack Overflow combined.

If you're not sure where to start on this, there's one area of financial literacy that is absolutely essential to understand, and that's tax.

Lesson 4: Learn the Tax System

There's a very famous clip of Kerry Packer (for many years, Australia's richest person), who was questioned about his tax practices in court back in '91. This is worth a quick watch (it's 2 minutes):



The key sentence being the last one in that clip:

Now, of course I am minimising my tax and if anybody in this country doesn't minimise their tax, they want their heads read because as a government, I can tell you you're not spending it that well that we should be donating extra.

Regardless of what you may think of the tax practices of billionaires, it's hard to argue with that statement (it's also hard not to chuckle just a little!). Tax is bloody complicated stuff yet it's something we all need to deal with in one way or another. It also consumes a significant chunk of your income and that only increases as you earn more and spend more. Understanding how your local tax system works is an absolutely essential part of that financial literacy I was just writing about.

For example, in Australia we have pretty attractive <u>negative gearing</u> tax laws for real estate and I'll steal the definition off Wikipedia to explain precisely what that means:

Negative gearing is a form of financial leverage whereby an investor

borrows money to acquire an income-producing investment and the gross income generated by the investment (at least in the short term) is less than the cost of owning and managing the investment, including depreciation and interest charged on the loan (but excluding capital repayments).

What this has meant for us is the ability to buy property and claim deductions for non-cash expenses (that is they're not actually coming out of your pocket) thus reducing our taxable income and ultimately increasing our take-home pay. For example, buildings, fittings and fixtures all "depreciate", that is their value decreases over time. Think about curtains - they wear out and need to be replaced and the Australian tax system affords you the ability to claim that depreciation before you actually need to spend the dollars. Your country may well have different laws, but the point is that tax constructs exist to help you legally reduce the amount payable. (Side note: there's been calls for years to abolish negative gearing in Australia in this fashion and there were indeed pretty significant changes made in the 80's... then rolled back.)

Retirement funds are another great example. In Australia, our "superannuation" scheme (think 401k in the US) makes it very attractive to contribute extra cash at a low tax rate. Only up to a threshold, that is, and even that changes based on your age but again, there are constructs designed by the government the help everyone maximise the effectiveness of the dollars they earn by minimising the amount of tax payable on them.

Tax is also where professional help is really important. Unless you're on a very low income that's just a simple wage from an employer, in my experience the ROI of professional guidance means it makes sense to get a good accountant early. Especially once there's more money involved, a very small percentage difference made by a taxation professional easily covers their cost (you may well find that's an allowable deduction too). Over time, our accountancy needs changed from a basic accountant we saw once a year to a larger scale firm we call on regularly. Your needs may well change too as you move through different phases of life, but get someone you can trust and get them early.

Optimising your tax position is free money. Free legal money and there are many, many ways to do it. In this industry, there's everything from income-producing equipment to conferences to <u>charitable donations to an organisation like Let's Encrypt</u> that can reduce your tax bill (obviously get expert advice on this if you're not sure). Sometimes, it's even just as simple as deferring tax that's payable so that you have access to the money for longer and can reap the benefits of the interest it earns. Pay your taxes, but don't donate extra.

Lesson 5: Know Good Debt from Bad Debt

The word "debt" immediately has negative connotations for a lot of people. Many of those people have a bunch of "bad" debt and little or no "good" debt. The latter term might sound paradoxical, but I'll get back to that. Let's start with the bad stuff.

Bad debt is the likes you have on a credit card. It's almost always accrued on a depreciating asset (for example, a new TV) and it's very often at a high interest rate. A credit card in Australia right now can easily run you around 20% per annum which means that not only are purchases going to cost way more than the sticker price (assuming the card isn't paid off each month), but the value of the purchase is also heading south leaving you with negative equity (you owe more than the thing is worth). Because credit cards have such a high rate on them, the single best investment you can make right now is almost certainly to pay off any debt you have on a card as fast as possible. Think back to that Vanguard chart - the highest yielding shares they had there (the US ones) were growing at 10.6% and paying off a credit card can effectively earn you double that. (Side note: that last sentence isn't entirely accurate as income earned on investments is usually subject to tax whereas paying off consumer debt will often have no tax obligations at all. Or if we go even deeper down the rabbit hole, those US shares at 10.6% include capital gains and that's something you may only pay tax on when you sell. So in other words, both the points made in

this side note make the investment value of paying off credit card debt even more important than investing in other asset classes.)

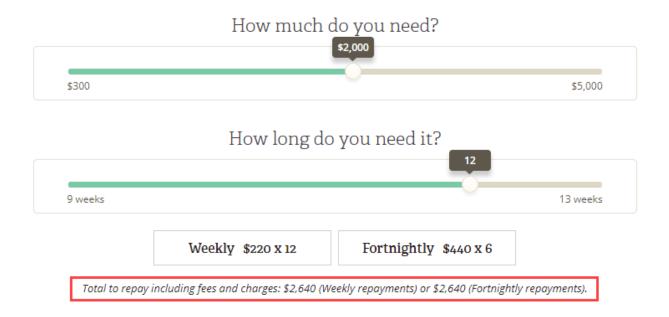
Payday loans are another prime example where you have fast, easy access to cash but pay an astronomically high interest rate for the privilege. For example, via <u>Nimble</u>, one of Australia's most prominent short-term lenders:

Small loans

If the principal amount you borrow is between \$300 - \$2,000 you'll pay:

- Establishment fee:
 20% of the principal amount
- Monthly loan fee:
 4% of the principal a month
- Payment options range from 62 days to 4 months. Actual loan term may vary depending on approved amount and your individual circumstances.

What does that look like in actual dollar terms? Let's imagine you need a couple of thousand for 12 weeks:



I highlighted the most important part in red because for some reason it was very small and a bit hard to read... In other words, for a loan that's less than a few months long you pay back an additional 32% over what you actually borrowed in the first place. This, in effect, makes whatever it is you bought with that money 32% more expensive and yes, I know that many people are under financial duress and may not have other options, the point is to understand what the actual impact of this debt really is. Remember also that compound interest works on debt too, not just savings. The longer you run with debt, the more you pay. (Side note: I watched a really interesting Netflix documentary on short term loans recently as part of their Dirty Money series - check out the Payday episode.)

Good debt is an investment. All our property purchases, for example, have loans not only because we simply couldn't have afforded to pay cash at the time, but because debt can give you leverage. Rather than paying, say, \$250k in cash, you'd put down, say, a 10% deposit and pay perhaps 5% per annum in interest. You then have cash flow from the asset (rent paid by tenants) and as mentioned earlier, there may also be non-cash deductions that give you taxation benefits. You also have expenses, primarily loan repayments but also maintenance, council rates, insurance and possibly strata and property management fees. I

don't want to go down that rabbit hole here (we're getting back to the importance of financial literacy again), but the point is that debt can be used to build wealth in an accelerated fashion. (Incidentally, the same approach can be used in shares and managed funds, this is not just the domain of real estate.) Borrowing for education can also be good debt. Kylie and I both had student loans via HECS in Australia which we had to pay off as we began earning money. This was an investment in our future and the return on the investment was an education.

This isn't to say that "good" debt is always a smart idea and in some cases, it can amplify losses dramatically. Some of the properties we bought were only several years old and were being sold for 30%+ less than the original owners had paid. Developers often entice buyers by offering "honeymoon" interest rates and rental guarantees that make the cash flow position look very positive to the unsophisticated investor. However, once those expired and the penny dropped that the properties were no longer financially sustainable (interest rates went up, rent went down), they became distressed sales and the unfortunate purchasers learned that the market valued the properties at a very different mark to what the developers were selling them for. Suddenly, the 5% deposit they paid to get access to real estate has created negative equity 6 times more than that.

Conversely, what we might traditionally consider "bad" debt can be good and I'll give you an example of that. This whole post kicked off with me talking about a car and as much as I love them, let me be really clear about this: fancy cars are one of the worst possible things you can ever spend your money on! They're functionally equivalent to models that are a fraction of the price, they depreciate very rapidly and they have a bunch of acquisition costs that disappear into thin air the moment you buy them (stamp duty, for example). But those are principles I understand very well so I make purchases with full consciousness of the financial impact. The point re bad debt potentially being good is that whilst a car is a depreciating asset, we've had cars in the past where the manufacturer's interest rate was far more attractive than the interest we could earn on the money elsewhere which would make paying cash a sub-optimal use of the

money. You need to be careful that an attractive interest rate isn't just capitalised into the purchase cost of the vehicle (and I've definitely seen that before), but the point is that debt can be used in a variety of constructive ways and some of them may be unexpected.

Over and over again, we come back to financial literacy and a big part of that is understanding not just how to use debt efficiently, but how to manage the risk it creates. I reckon as technical folks we tend to be more analytical than your average person and one of the best things you can do for you financial wellbeing is to chuck everything into spreadsheets. This debt situation, for example, can be really multifaceted so if you're looking at taking out a loan, put everything into Excel and analyse the bejesus out of it; cash flow impact, capital gain / loss, opportunity cost (what else you could do with the deposit and repayments), etc, etc. I've had many occasions in the past where I've literally sat down and written all my analysis in C# because I understood the code better than the finances! But by doing that, you learn, and that's a great way of working on financial literacy. (Side note: service like Mint are also a great way of tracking your financial position.)

Lesson 6: Diversify Earning Potential and Risk

This one starts to get to the heart of where money comes from and how to protect it. Specifically to this industry, we have much better potential than most to both earn it and keep it - let me explain.

Traditional incomes generally boil down to trading time for money from a single source (your employer). It certainly did for me for many years and in my case, it meant going into Pfizer each day, doing my architecty thing and receiving a monthly pay check. As we've already established further up, a software architect in this industry can do quite well but this traditional means of working does

create risk and I saw that manifest itself through many rounds of redundancies over the 14 years I was there. I'd see the stress people went through as their roles were cut and they were out of a job and there were 2 main reasons for that:

- 1. Their job was their sole source of income and if it went, so did their cash flow (in some cases, it was the sole source of income between both them and their partner who may be a stay at home parent)
- 2. They were worried about their ability to get another job which, again, would also have a pretty significant cash flow impact

We went through that stress ourselves; about 7 years ago Kylie's job was made redundant. She was 6 months pregnant (seriously, who does that to a pregnant woman about to go on maternity leave?!) and it was entirely unexpected and left us with a very uncertain future. Fortunately, we had my income to cover us and we'd obviously planned in advance for the maternity leave, but it still rocked us.

So let's drill down on this "diversify earning potential" concept and the first point I want to make on that is about your own personal marketability. My very first blog post ever was Why online identities are smart career moves and a cornerstone of that post was that you never know when you might be looking for another job. Making yourself marketable isn't something you can do well at the drop of a hat, it can take significant effort and it's something you need to plan for in advance. You might not necessarily think of that as a personal finance tip, but it can have a fundamental impact on your ability to earn money.

One of the suggestions I received when tweeting about this post last week was this:

How to get rich and not get pwned! Yes please. Maybe some advice about investing in your own learning and pet projects would be nice.

— ₱ Ruan Kranz? (@krankit_io) December 27, 2018

Here's a perfect example that illustrates my point: when I first interviewed for the Pfizer job in 2001, I showed a pet project I'd built. It was a classic ASP and Access Database (stop laughing) project that managed photos I'd taken. It was very basic, but it gave me something to show that demonstrated work in the field. I clearly remember showing my boss's boss the work and him being impressed by it, despite its simplicity. This was a personal project done on my own time as part of my own education and it played an important part of landing me the job I had for the next 14 years. That's the job that contributed significantly to the investment portfolio!

Pet projects, open source contributions, robust Stack Overflow profiles, local user group engagements and a raft of other things you can do in your spare time all contribute to marketability and in turn, diversify your earning potential. (Incidentally, the talk I referenced earlier on Hack Your Career covers all of this in more detail.) This is one of the great advantages we have in this industry in that it's so easy to expand our professional repertoire in our spare time. I'll give you an example of the antithesis of that: One of the people I saw forced into redundancy at Pfizer was in a senior role they'd been in for a very long time (I'm going to be a little vague here in case they read this) and frankly, they really had very little (any?) industry experience outside of that. They were proficient at their job but they really didn't have skills that were transferable across the industry and when the redundancy finally came, they were out of work. Permanently. They ended up re-skilling in another industry in what was quite a stressful time for them.

Moving on, another great attribute of tech is the ability to diversify income sources. Now, I'm conscious there are cases where the employer may prohibit some of these things (even on personal time) but as an example, I did a lot of small independent website projects whilst in my corporate role. Nights, weekends, holidays were often spent building brochureware websites or other little pieces of work that could earn income. Independent income which would contribute to our financial wellbeing. That money then went into the property portfolio and grew further so think about the leverage that provided: the extra

money earned was nice in and of itself, but that was then used to borrow money (so it was leverage) which bought appreciating assets. Those nights, weekends and holidays ultimately became very valuable.

In 2012, I started creating what many people came to know me by: Pluralsight courses. Again, something that could be done independently without conflict with my day job. There are many, many little opportunities like this which can actually contribute to both the points made in these last 2 paras, namely diversifying your experience and actually generating income. Today, no more than about 20% of our income comes from any one source which is enormously important in terms of diversification. It means that, for example, if Pluralsight goes down the toilet then yes, I'd be very upset by that but no, it wouldn't be a life-changing event.

Which brings us to risk. Risk is reduced when you have more choices and it's reduced again when you have more sources of income. Drawing it back to investment strategies, you'd never proverbially put all your eggs in one basket by, say, putting all your cash into one stock. Using all your savings to buy that one magic bean, so to speak. When we bought real estate, we bought at a level that would enable us to diversify; I'd rather have 2 small apartments in different suburbs than 1 house because it gives you insurance against everything from tenant vacancies to repairs that need to be made to something extreme like the place burning down. And you definitely don't want all your exposure in one asset class either; property, shares, cash and all sorts of other investment vehicles enable you to spread risk. Try a Google search for life savings lost in investment scheme and you'll understand why this is so important.

Invest in diversifying your earning potential and your assets such that it reduces your risk.

Lesson 7: Prepare for Luck

When I started drafting this blog post all those years ago, one of the things I immediately thought of was this book:



Outliers



THE STORY OF SUCCESS

Malcolm Gladwell

#1 bestselling author of The Tipping Point and Blink

Malcolm Gladwell is a sensational author and his previous books <u>The Tipping</u> <u>Point</u> and particularly <u>Blink</u> are absolute must reads. But what I particularly liked about Outliers is how he systematically broke down the factors that contributed to the success of very noteworthy people such as Bill Gates, The Beatles and even elite athletes. On that final point, let me draw an extract from

Wikipedia that illustrates one of the success factors Gladwell identified:

The book begins with the observation that a disproportionate number of elite Canadian hockey players are born in the earlier months of the calendar year. The reason behind this is that since youth hockey leagues determine eligibility by calendar year, children born on January 1 play in the same league as those born on December 31 in the same year. Because children born earlier in the year are statistically larger and more physically mature than their younger competitors, and they are often identified as better athletes, this leads to extra coaching and a higher likelihood of being selected for elite hockey leagues.

There are 2 ways of thinking about this as it relates to success factors and the first is that elite hockey players are exceptionally talented. Regardless of the other opportunities that were granted to them, you simply can't play at that level unless you're at the absolute top of your game. The second is the real insight in this piece and it's that the older kids have a natural advantage due to those extra months of growth. An unfair advantage, some would argue, but an advantage all the same. But it wasn't all luck either - there's plenty of kids born in January that can't compete with much younger players because they simply don't have the natural talent or the family support or the dedication to train or whatever else it may be. Being successful at that level requires both luck and talent.

Bill Gates is worth a mention as it ties in nicely with the tech-centric theme of this post. Yes, he's obviously a super smart bloke, but it was his (very fortunate) access to computers courtesy of his mother's job that amplified that talent and enabled him to build Microsoft. And this is really the point I'm getting at in this lesson: we all come across fortuitous situations - "luck", if you will - and you need to be prepared to take advantage of those opportunities. Those situations may be anything from a sudden job opportunity to a chance investment, both of which often require preparedness to take advantage of. For example, do you have a presentable resume and references for that chance job? Do you have up to date tax returns and financial statements for the investment? Are you able to leverage

the skills and the assets that you have - that we all have - to be able to take advantage of these opportunities when they arise? I certainly haven't always and I lament the ones I missed because I simply wasn't prepared.

I vehemently dislike seeing successful people referred to as "lucky" or "fortunate" without further context. Not because they're inherently wrong words to use, but because they imply people achieved that success by chance. It must also be disheartening for others who don't believe they're as lucky or as fortunate themselves which is why I love this quote:

I am a great believer in luck. The harder I work, the more of it I seem to have.

There's debate about who originally said it but it doesn't particularly matter as the sentiment rings true regardless. What I hope people take away from it is an acknowledgement that hard work and preparation amplifies the luck that we all come across from time to time.

One more thing on the whole "luck" piece because it will come through in comments if I don't address it: Just as the older hockey players benefited from the month they were born in, I've benefited from factors I was born into. My gender. My ethnicity. The country I was born in. Even the countries I've lived in; I spent the last few years of high school in Singapore which was an absolute tech mega centre compared to most of the rest of the world in the early to mid-90's when I was there. A chance meeting at the local windsurfing club with a guy working for a satellite systems engineering company in '92 got me my first part time job in technology. These are factors I had no control over, but I amplified that good fortune by working my butt off when I was given the chance. Whatever your circumstances, the premise that opportunities will present themselves over time and that being prepared to leverage them is important is still an absolutely essential lesson.

Lesson 8: Put a Price on Your Time - and Your Family

I stopped playing video games probably about a decade ago. Half Life 2 was my game of choice at the time and I could easily blow a few hours fragging everything that moved. Whilst it certainly wasn't at an addiction level, it was still enough time spent that eventually it dawned on me that it simply wasn't a good way to invest my hours. Now I want to be clear about something here too: investments aren't always of a monetary nature, they can be investments in your health or your mental state or your family and as it stands today, I spend more time playing tennis each week than I did fragging. But the return on that investment is so much greater for my mental state and my health than what HL2 ever was.

To the point about putting a price on your time, I realised holistically I was much better off focusing on our investments and my own personal development than I was spending the time gaming. As time has gone by, I've become more and more conscious of what the value of my time is. Sometimes it's a clear monetary value; I charge companies to run <u>security workshops</u> which is a direct exchange of time and money. Other times it's much less tangible but it feels like it's moving things along in the direction I want them to go. This blog post is a perfect example of that insofar as it will make me zero dollars directly but I feel like it's the right thing to do because it has the potential to improve life for others. Understanding the value of time (and particularly how it changes over the years) has also helped me decide where to spend money to buy back hours; a house cleaner a couple of times a week, someone to wash the cars, business class airfares.

Then there's putting a price on your family. People hate it when I use this term - "what do you mean I should put a price on my family, my family is priceless!" - and they continue to hold that position as they head off to the office each day. The reality is that we all trade time with our families to partake in activities that

enable us to actually support them, but most people don't favour thinking about it in those terms. It doesn't have quite the same ring to it, but perhaps a more accurate title would be "consider how much family time you're willing to sacrifice for your prosperity and how long you're willing to wait for that investment to pay off". If you don't have a family to consider, put it in terms of other personal activities you're willing to trade; I traded gaming, others might trade social activities or a holiday or some other form of sacrifice that results in them working towards their own prosperity.

What I mean by putting a price on your family is that you should work out when it makes sense to prioritise spending time with them and when it makes sense to invest time to focus on other things. For many people, there's no desire to commit anything more than 40 hours a week to earning a living and that's just fine, so long as the lifestyle that gives them is consistent with the one they want and they're not left unfulfilled as a result. What drives me nuts is when you see people wistfully longing for certain financial or lifestyle goals yet being unwilling to make the sacrifices to get there (more on that in the next lesson).

My balance has changed over time. In the earlier years of my career when I was mostly on hourly contracts, it would be 11-hour days most of the time because surprise, surprise, that pays a lot more than 8-hour days (and remember, that went into leveraged assets that then grew in value over many years). It was fine when it was before kids too and Kylie was either studying or building her own career with similar hours, we both just knuckled down and got on with it. It's always going to be harder with kids, particularly because higher workloads are inevitably passed onto your spouse if you're the one doing the extra career things. What I'm finding now is that because we made those sacrifices before kids were around, we're enjoying the pay-off while they're still young.

If nothing else, at least consciously make choices about where time is spent and one of the best things that'll help you do that is to have a goal.

Lesson 9: Have a Goal

The best way I can explain this is to share a speech by Arnold Schwarzenegger. Invest 12 minutes listening to this:



Don't waste your minutes. Work your arse off. You have 24 hours in a day, you sleep 6 of them, maybe you burn 12 with work and travel so now you have 6 hours left. You eat / schmooze a little, but you see how much time is left.

If you don't have a vision of where you're going, if you don't have a goal where you go, you drift around and you never end up anywhere.

A goal keeps you focused. A goal drives you to invest time in working towards something. A goal makes you relish the pain required to achieve it. Schwarzenegger talks about the physical pain of reaching his goals, but also about improving knowledge by investing time which aligns with what you've read here. Now, clearly he took a very extreme approach to reaching his goal because it was an extreme goal. I'm not saying everyone should go out and spend every spare moment figuring out how to maximise their dollars, but what

I am saying is that you need to know why you're doing this - what you're working towards - and depending on how lofty that goal is, it may indeed take a significant amount of effort over a long period of time.

In this industry, we work with goals the whole time and we've all worked with tools that help enable us to hit them. We have backlog items that need to be completed and they can just as well be things like getting your insurances in order, assessing your retirement strategy (yes, even when you're young) or setting a learning objective. We deliver work units in sprints and when you have a long-term goal, there's going to be many individual sprints within it. Kylie and I continually have retrospectives; what's working, what's not, what do we need to do differently. And if we really want to draw out the agile analogies, nothing requires adaptive planning like your financial future does because there are so many environmental factors that change; your job, your family structure, interest rates and any number of other things that require a course correction. We, tech friends, understand this. This is what we do day in and day out and you can extend that to your personal financial prosperity.

Inevitably, we all have multiple goals and they'll change over time too; for many years whilst I was living in Sydney and working for Pfizer, my goal was to gain independence and move back the Gold Coast where my family was. In 2015, we did that:





Just done the big move with @KylieMHunt, totally stoked about our new house on the Gold Coast!



11:23 PM · Oct 15, 2015



So, I made new goals. I've certainly had others too and they haven't always been this long-term or life-changing. For example, I've had goals for certain cars I've wanted and in some cases, it's taken many years to achieve them. In other cases, I'm yet to achieve them but they're still there on the horizon, driving me forward and giving me direction.

Goals can be very personal; perhaps your goal is to retire young. Maybe it's to support your extended family. It might even be to give as much as you can to charity (Gates is a perfect example of that) and all of those are just fine, but have a goal because without that... you drift.

Lesson 10: Financial Prosperity is a Partnership

I wanted to finish on this point because it's absolutely pivotal to making all the previous ones actually work. If you're in a partnership with someone (wife, boyfriend, whatever), perhaps the most valuable advice I can give is that you must approach financial prosperity as a partnership with a shared vision. If you're not aligned - if you have fundamentally different objectives - you won't be able to give your goals the focus they deserve.

I think back to friends I've seen in the past struggle with this. For example, one partner becomes resentful of the amount the other is spending on personal indulgences. Or they resent the family sacrifices the other is making. Or one is satisfied with a subsistence living whilst the other dreams of millions. Lack of alignment not only makes achieving financial objectives difficult, it can drive a wedge right through the middle of a relationship and I'm sure we've all seen many fail simply because the couple don't see eye to eye on fundamental issues.

I'll give you a few examples of what I mean and the first one that came to mind (for some strange reason) was when Kylie and I were planning a family. Like most couples, there comes a time where that's on the cards and for us we started talking seriously about it in 2008. As we began planning, we literally went to a quiet spot in a local restaurant with a laptop and drew up a spreadsheet of what having a baby would mean. We did this together and planned everything from loss of income due to maternity leave, government parental benefits, the taxation implications of both and even medical expenses and the maintenance cost of a child. I'm sure we didn't get that all spot on (the last one in particular), but the point is that we made a financial decision together (and having a kid is a very big financial commitment) with as many of the facts as possible in front of us.

That partnership extends to everything from the investments we make to the travel I do to the insurances we have (NB: things like income protection and life

insurance are another one of those financial literacy things). This isn't just to ensure alignment, it's also a great sanity check. If you're in a relationship, you'll probably find there are aspects of this whole financial prosperity thing that each of you does better than the other; I'm "big picture" and number orientated, Kylie is detail-focused and frankly, much more patient than me! Explaining things to each other has a way of ensuring you stay on track.

But perhaps even more importantly than all of that, relationships are meant to be a partnership. A journey you take together. Hopefully a very long journey that requires planning and there are few more fundamental relationship issues than how you view money.

Summary

If you're working in tech, you're working in one of the most well-paid industries with the greatest growth potential and career prospects out there. Your financial potential almost certainly exceeds that of almost everyone else around you. You're already winning just by being here and my hope is that whilst the first tweet in this post might have provided motivation, the post itself helps provide inspiration.

```
Nice wheels, mate.

Motivation No. 946,624 to start hacking my career this year. :)

Keep up the good work, you've earned it all and then some.

— BlueTeam_Ninja - Side Hustling (@BlueTeam_Ninja) December 26, 2018
```

Feel free to ask questions in the comments section below and I'll answer what I can. Also - and I trust this was obvious already - do treat this post as a reflection of my own views and experiences and get professional advice where necessary.

Comments

Appreciate you sharing your experience and knowledge. I wish I'd applied the same focus and consideration from an earlier age :) One thing I'm curious about - when you became an independent consultant, did you give a lot of consideration to specialising in Infosec, or was it something you already had an interest and passion for? I'm guessing your background in corporate IT would have given you a broad knowledge base, so I'm wondering if targeting Infosec was part of the master plan. If so then great choice and insight: I'm sure it's a lot more lucrative than specialising in, say, databases or networking for example.

Troy: This post helps answer your question: https://www.troyhunt.com/to...

I worked on contract until about 2007 when I went permanent with Pfizer. First for years as an independent then for years contracting to them, all as a developer. When they made my role redundant 4 years ago, I already had Pluralsight earning twice as much money as Pfizer so there wasn't much to need to think about! (More on that in the Hack Your Career talk.)

So in short, the infosec focus was organic and not something I consciously set out to do in lieu of my previous corporate architect role.

__

As someone with a finance background, this is awesome advice, even for those outside the tech industry.

I feel like adding something to what you said about taxes however, which is to not only be aware of the local tax code, but to really understand the cost the deductions you're looking to take.

Case in point: In the US, you can take a deduction on your federal taxes for the interest paid on the mortgage of your primary residence. This is pretty nice, as it (at least slightly) minimizes the barriers to new home ownership.

However, I know of people who would pay off their house in full, then turn around and take out another mortgage just so they can claim that deduction. This makes no sense from a

financial perspective, as you're paying 3x as much in interest just to avoid paying some taxes. If the same person would instead invest that money, they would pay a little more in federal taxes but they would be not only keeping the money that would have been spent, but could be earning interest on it themselves. Even if they didn't invest it but stashed it under their bed, they would still end up with a net increase.

In the end it all goes back to financial literacy, which I agree is sorely lacking these days. I'm a big fan of Dave Ramsey, who teaches how a person can increase their net worth, and while I don't necessarily agree with him on a few things, his logic, reasoning and advice are sound and can be applied to pretty much anyone.

Troy: This is reminiscent of people in Australia seeking out investment properties running at a loss solely to reduce their tax burden. Want to reduce your tax burden? Pay more interest or charge less rent! I suspect based on your comment as well that this is a problem with not being able to holistically look at the cost of financial decisions. Your comment also raises the issue of the opportunity cost incurred by committing money to one pursuit when it could be used for another which seems like good material for a follow-up post $\ensuremath{\mathbb{C}}$

Epilogue

What really hits me reading back over this post is points 1 and 10, namely how money buys choices and then how financial prosperity is a partnership. These bookends have been particularly poignant to look back on whilst writing this book, let's start with the first one:

I wrote about the positive choices that money can give you, for example I like nice cars and without sufficient resources that's simply a choice I wouldn't have. But with the benefit of time, I also saw how money can empower people to make choices that may ultimately be detrimental personally or in our case, to our relationship. Having the choice for my wife at the time not to work was slowly eroding the shared vision we'd once shared, especially as the demands

on me increased. Aspects of having that choice were good, for example her being able to focus on her health. But aspects of that choice were also detrimental; the decision to leave the workforce significantly predated the aforementioned health issues and clearly, her priorities in life were changing. As work was left behind, so many of the benefits of work also disappeared. Social connections. A sense of purpose and fulfilment. A stable daily routine. And, of course, financial contribution to the family. These were all hurting our relationship and ultimately, significantly contributed to its demise. Having money made not working easy in the worst possible way.

When I wrote about financial prosperity being a partnership, I was spot on, but it was only when writing this epilogue that it really hit me: this section of the blog post was what I wanted it to be, not what it was by that time, and it was ultimately lack of alignment that became a big issue for us. To use my own words, lack of alignment "drove a wedge right through the middle of the relationship" not so much as it related to the goals (she certainly still wanted the lifestyle we'd built), but rather the contributions required from each party to reach them. Having a lack of alignment isn't about laying blame in one direction or the other, rather it's about recognising that if you're not on the same wavelength about something this fundamental, sooner or later you're going to have some serious issues to deal with. But when we ultimately did have to deal with those issues, money gave us choices. The choice to continue living well despite splitting assets. The choice for her to decide if and when she wants to work. The choice to keep the kids in private schools and all sorts of other choices so many people going through a divorce simply don't have. I can't imagine how much more difficult it would have been without money.

So much more of this post also stood the test of time, for example how essential it is to diversify earning potential and risk. None of us saw COVID coming, but that's precisely the sort of unexpected outcome you have to plan for, namely that you might suddenly have an income source dry up. Or both income sources dry up. The job market might tank. You might not even be allowed out of your house! But COVID also created opportunity for many people which brings me to the point on preparing for luck: I was speaking to

the guy who cleans our pool during the height of the pandemic, and he was raking it in - "everyone is staying at home now and they're in the pool all the time, so business is great!" He was able to capitalise on the opportunity this presented to his business because he was in the right place at the right time, just purely by luck.

Lastly, when preparing for this book and asking people what they'd like to read, someone enquired about how I felt with my decision to travel a lot which is addressed in the point about putting a price on your family. So, how do I feel? Absolutely fine. I feel fine because what I did was shuffle the time I spent with my family around rather than just deduct from it. I'd travel for extended periods but then when I was home, I'd be at every single tennis match my son played and I'd almost always be the one to drive the kids to and from school or take them out on the jet ski. Plus, the nature of my life gave me heaps of opportunities for the kids to travel as they joined me on trips to the US, Canada, Hawaii, Norway, the UK and just for fun, taking time out each year to go skiing in Australia too. The quality of the time I had with them was without doubt greater than when I worked a "normal" job in my previous life. These days, it's even more so as they divide their time evenly between 2 households; I have a week where we squeeze out every minute of every day doing something awesome together, then I throw myself into work for a week and put all my focus into that. Frankly, it's the best work and family balance I've ever had and were it not for the earlier sacrifices and access to money I have now, it'd be a much bleaker story.

HOW TO TRACK YOUR KIDS (AND OTHER PEOPLE'S KIDS) WITH THE TICTOCTRACK WATCH

I feel that over the years of seeing so many security vulnerabilities and data breaches I'm at the point now where I look into something and at a glance, have a pretty reliable sense of "yeah, this thing is gonna be terrible". I read Malcolm Gladwell's book "Blink" years ago ("The Power of Thinking Without Thinking") which firstly, is an absolutely sensational book and secondly, made a lasting impression on me in terms of reflecting on how I make decisions quickly. When Ken Munroe contacted me about this watch, I had that Blink moment and started preparing myself for the impending train wreck that would be the security posture of this watch. And it was. And then it was again later on, but I'll save that for the epilogue.

15 APRIL 2019

o you ever hear those stories from your parents along the lines of "when I was young..." and then there's a tale of how risky life was back then compared to today. You know, stuff like having to walk themselves to school without adult supervision, crazy stuff like that which we somehow seem to worry much more about today than what we did then. Never mind that <u>far less kids go missing today than 20 years</u> ago and <u>there's much less chance of them being hit by a car</u>, circumstances are such today that <u>parents are more paranoid than ever</u>.

The solution? Track your kids' movements, which brings us to TicTocTrack and

the best way to understand their value proposition is via this news piece from a few years ago:



Irrespective of what I now know about the product and what you're about to read here, this sets off alarm bells for me. I've been involved with a bunch of really poorly implemented "Internet of Things" things in the past that presented serious privacy risks to those who used them. For example, there was <u>VTech back in 2015</u> who leaked millions of kids' info after they registered with "smart" tablets. Then there was <u>CloudPets leaking kids voices</u> because the "smart" teddy bears that recorded them (yep, that's right) then stored those recordings in a publicly facing database with no password. Not to mention the various spyware apps often installed on kids' phones to track them which then subsequently leak their data all over the internet. <u>mSpy leaked data</u>. <u>SpyFone leaked data</u>. <u>Mobiispy leaked data</u>. And that's just a small slice of them.

And then there's kids' smart watches themselves. A couple of years back, <u>the Norwegian Consumer Council discovered a whole raft of security flaws in a number of them</u> which covered products from Gator, GPS for barn and Xplora:



These flaws included the ability for "a stranger [to] take control of the watch and track, eavesdrop on and communicate with the child" and "make it look like the child is somewhere it is not". These issues (among others), led the council's Director of Digital Policy to conclude that:

These watches have no place on a shop's shelf, let alone on a child's wrist.

Referencing that report, <u>US Consumer groups drew a similar conclusion</u>:

US consumer groups are now warning parents not to buy the devices

The manufacturers fixed the identified flaws... kind of. <u>Two months later, critical</u> <u>security flaws still remained in some of the watches tested</u>, the most egregious of which was with Gator's product:

Adding to the severity of the issues, Gator Norge gave the customers of the Gator2 watches a new Gator3 watch as compensation. The Gator3 watch turned out to have even more serious security flaws, storing parents and kids' voice messages on an openly available webserver.

Around a similar time, <u>Germany outright banned this class of watch</u>. The by-line in that piece says it all:

German parents are being told to destroy smartwatches they have bought for their children after the country's telecoms regulator put a blanket ban in place to prevent sale of the devices, amid growing privacy concerns.

Wow - destroy them! The story goes on to refer to the German Federal Network Agency's rationale which includes the fact that "parents can use such children's watches to listen unnoticed to the child's environment". This is a really important "feature" to understand: these devices aren't just about tracking the kids whereabouts, they're also designed to listen to their surroundings... including their voices. Now on the one hand you might say "well, parents have a right to do that". Maybe so, maybe not, you'll hear vehement arguments on that both ways. But what if a stranger had that ability - how would you feel about that? We'll come back to that later.

Around a year later, <u>Pen Test Partners in the UK found more security</u> <u>bugs</u>. Really bad ones:

Guess what: a train wreck. Anyone could access the entire database, including real time child location, name, parents details etc.

This wasn't just bad in terms of the nature of the exposed data, it was also bad in terms of the ease with which it was accessed:

User[Grade] stands out in there. I changed the value to 2 and nothing happened, BUT change it to 0 and you get **platform admin**.

So change a number in the request and you become God. This is something which is easily discovered in minutes either by a legitimate tester within the organisation building the software (which obviously didn't happen) or... by someone with malicious intent. The Pen Test Partners piece concludes:

We keep seeing issues on cheap Chinese GPS watches, ranging from simple Insecure Direct Object Request (IDOR), to this even simpler full platform

take over with a simple request parameter change.

Keep that exploit in mind - insecure direct object references are as simple as taking a URL like this:

example.com/get-kids-location?kid-id=27

And changing it to this:

example.com/get-kids-location?kid-id=28

The level of sophistication required to exploit an IDOR vulnerability boils down to being able to count. That was in January this year, fast forward a few months and <u>Ken Munro</u> from Pen Test Partners contacts me. He's found more serious vulnerabilities with the services these devices use and in particular, with TicTocTrack's product. He believes the same insecure direct object reference issues are plaguing the Aussie service and they needs someone on the ground here to help establish the legitimacy of the findings.

To test Pen Test Partners' theory, I decided to play your typical parent in terms of the buying and setup process and use my 6-year old daughter, Elle, as the typical child. She's smack bang in the demographic of who the watch is designed for and I was happy to give Ken access to her movements for the purposes of his research. So it's off to <u>tictoctrack.com.au</u> where the site leans on its Aussie origins:



I can understand why companies emphasise the "we host your data near you" mantra, but in practical terms it makes no difference whether it's in Australia or, say, the US. You're also often talking about services that are written and / or managed by offshore companies anyway so where the data physically sits really is inconsequential (note: this is assuming no regulatory obligations around colocating data in the country of origin). The "we take the security of your data seriously" bit, however, always worries me and as you'll see shortly, that concern is warranted.

The Aussie angle comes up again further down the page too:

Monitor from anywhere

Our Australian designed software platform allows you to monitor your child from anywhere in the world. Overseas excursions can be stressful but with TicTocTrack® you can still keep an eye on them from the comfort of your home.

At this point it's probably worthwhile pointing out that despite the Aussieness

asserted on the front page, the origin of the watch isn't exactly very Australian. In fact, the watch should be rather familiar by now:



So for all the talk of TicTocTrack, the hardware itself is actually Gator. In fact, you can see exactly the same devices over on <u>the Gator website</u>:



It's not clear how they arrived at the conclusion of "the world's most reputable GPS watch for kids and elders", especially given the earlier findings. And who is Gator? They're a Chinese company located in Shenzhen:



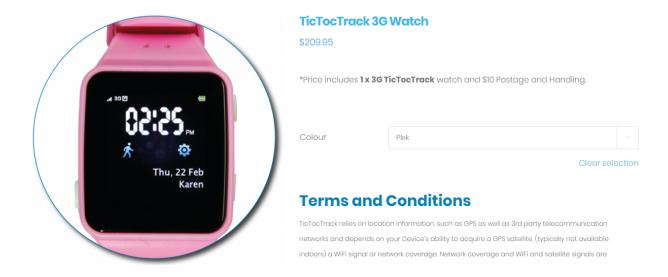
High-quality GPS watches are guaranteed

Gator Group Co., Ltd is located in **Shenzhen** which is one of the 5 international largest and wealthiest cities of **China**.

Gator's factory is **ISO 9000** certified which is a set of international standards on quality management to ensure the best GPS Watches produced effectively.

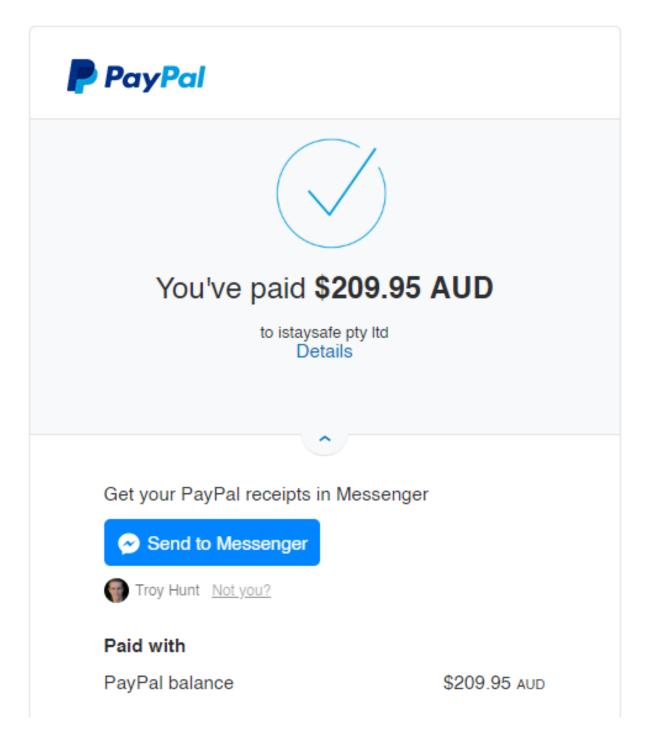
All in all, our mission is to offer the high quality and safety products, so we also got other important certificates, such as CE, SAR, FCC, CTL, RoHS, C-TICK LVD, etc. The country of origin would be largely inconsequential were it not for TicTocTrack's insistence on playing the Aussie card earlier on. It's also relevant in light of the embedded media piece at the start of this blog post: this isn't "a new device developed by a Brisbane mother" nor is the mother "the creator of the watch". In fairness to Karen Cantwell, it wasn't her making those claims in the story and the media does have a way of spinning things, but it's important to be clear about this given how this story unfolds from here.

Regardless, let's proceed and actually buy the thing. I get Elle involved and allow her to choose the colour, with rather predictable results:



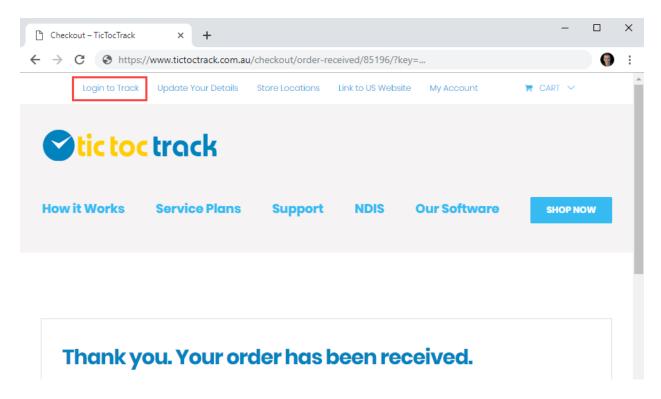
The terms and conditions were actually pretty light (kudos for that!) but <u>the link</u> to the privacy and security policies was dead. I go through the checkout process and buy the watch:

istaysafe pty ltd

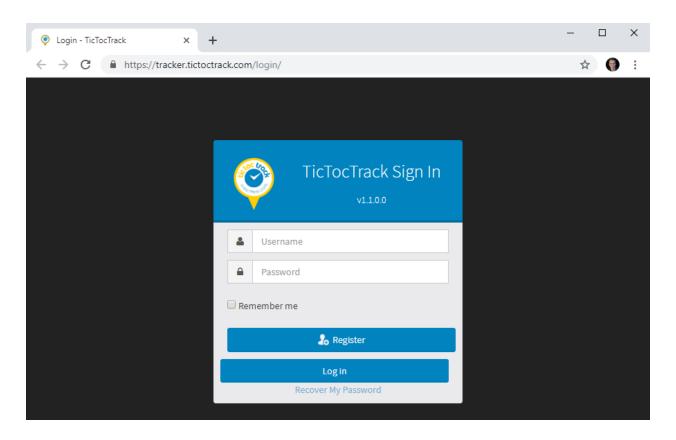


<u>iStaySafe</u> Pty Ltd is the parent company and we'll see that name pop up again later on. An email promptly arrives with a receipt and a notice about the order

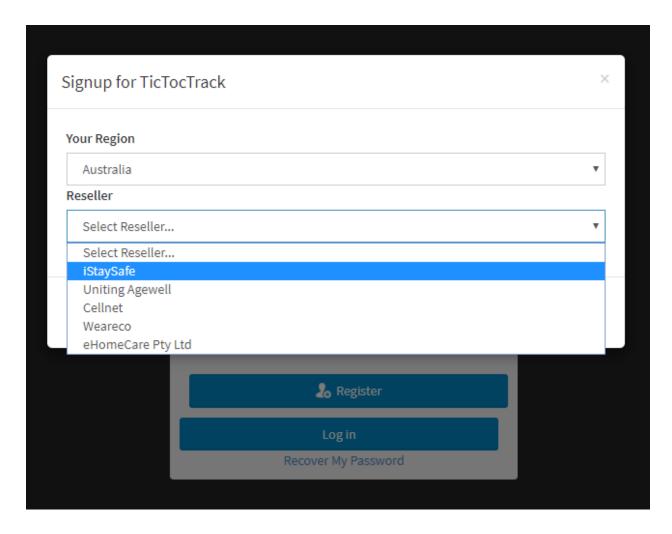
being processed, albeit without a delivery time frame mentioned. With time to kill, I decide to poke around and take a look at how the tracking works, starting with the link below:



Turns out the tracking app is a totally different website running on a totally different hosting provider in a totally different state:



<u>The primary site is down in Melbourne</u> whilst <u>the tracking site is in Brisbane</u> per the info on the front page. My credentials from the primary site don't work there and registering results in me needing to choose a reseller:



Here we see iStaySafe again, but it's the other resellers (all Aussie companies) that help put the whole Gator situation in context. <u>Uniting Agewell</u> provides services to the elderly and when considering the nature of the Gator watch, it made me think back to a comment on the Chinese manufacturer's website: "the world's most reputable GPS watch for kids and elders". <u>Cellnet</u> is a publicly listed company <u>with a heap of different brands</u>. <u>Weareco</u> produces uniforms. <u>eHomeCare</u> provides "smart care technology for healthy ageing" and <u>their product page on the GPS tracking watch</u> explains the relationship:



GPS TRACKING

GIVE YOUR LOVED ONES THE FREEDOM TO EXPLORE



GPS tracking watch to keep your loved ones safe. Complete with an active SIM all set up and ready to go, TicTocTrack@ allows you to locate your loved one from your smart phone or computer.

Keep in touch throughout the day. Whether it's a call to let them know you are running late or just a check to see they arrived or returned home safely, TicTocTrack® can give you peace of mind.

As it turns out, attempting to sign up just boots me back to the TicTocTrack website so I assume I just need to wait for the watch to arrive before going any further. Still, this has been a useful exercise to understand not just how the various entities relate to each other, but also because it shows that the scope of this issue isn't just constrained to kids, it affects the elderly too.

A few days later, this lands in the mail:





I'm surprised by how chunky it is - this is a big unit! For context, here it is next to my series 4 Apple Watch (44mm - the big one):

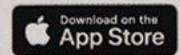


I'm not exactly expecting Apple build quality here (and as you can see from the pic, it's a long way from that), but this is a lot to put on a little kid's wrist. You can see the access port for the physical SIM card (more on that later), as opposed to Apple's eSIM implementation so it's obviously going to consume a bunch of space when you're building a physical caddy into the design to hold a chip on a card.

Regardless, let's get on with the setup process and I'm going to be your average everyday parent and just follow the instructions:

Quick Start Guide

Step One Begin by downloading the TicTocTrack® app. At the login screen choose 'Register' to create your user account.





Step

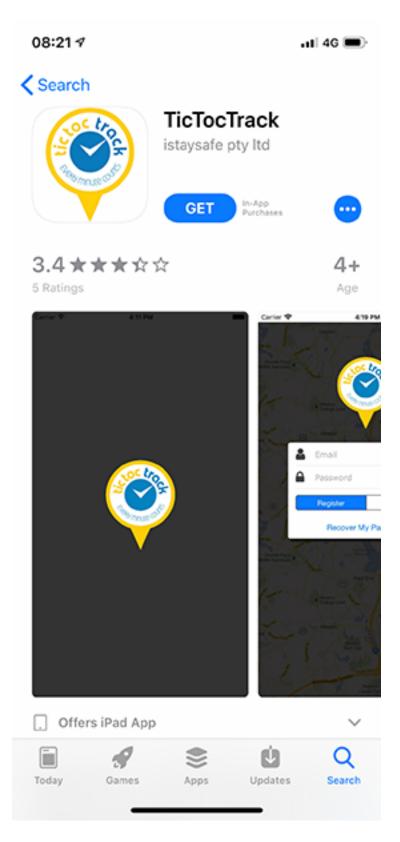
Follow the instructions to add your watch and subscribe to a service plan. If you are using your own SIM please ensure you have an active Nano SIM with voice, data and SMS credit. You will also need to know the APN name for your SIM. We recommend the Telstra network to ensure the best coverage. Alternatively, you can use the Telstra SIM supplied in your packaging and choose the 'Full service plan' option.



Once you have completed the process of adding your watch all you need to do is charge it and pop it outside for 15 minutes to generate a location.

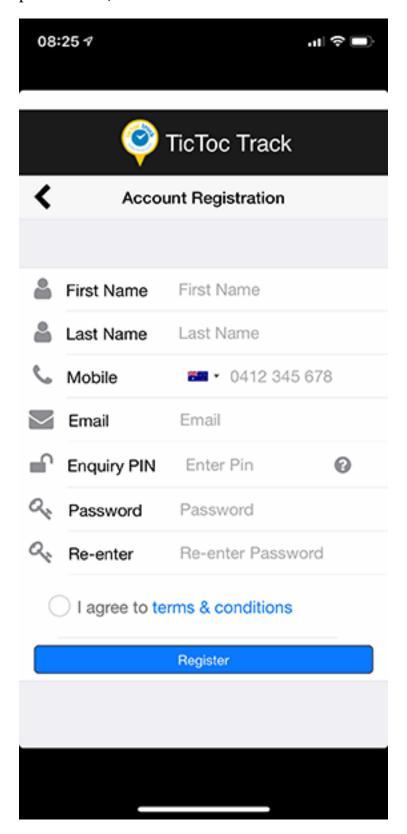
Turn over for some frequently asked questions.

The app is branded TicTocTrack and is published by iStaySafe:



Popping it open, the first step is registration (the mobile number is a pre-filled

placeholder):



I'm surprised by the empty space at the top and the bottom - just which generation of iPhone was this designed for? Certainly not the current gen XS, does that resolution put it back in about the iPhone 5 era from 2012? That'd be iOS 6 days which their <u>user manual</u> seems to suggest:

Support Platforms

- iOS (6.0 +)
- Android (4.2 +)





Supported devices

iOS

Compatible with iPhone, iPad.

Android

Support more than 8006 + android devices.

Whilst the aesthetics of the app might seem inconsequential, I've always found that it's a good indicator of overall quality and is often accompanied by shortcomings of a more serious nature. It's the little things that keep popping up, for example the language and grammar in the aforementioned user manual. Why is it "Support Platforms" and then "Supported devices"? And why is the opening sentence of the doc so... odd?

Welcome to TicTocTrack® User Manual! You are about to begin your

journey with the live tracking with your family.

That sort of language appears every now and then, for example in the password reset section:

If you forget your password, please use web portal to obtain new password.

It has me wondering how much of this was outsourced overseas and again, that wouldn't normally be worth mentioning were it not for the emphasis placed on the Aussie origins of the service (I know, despite it being a Chinese watch). The actual origins of the service become clear once you look at the download links for the app:

Download TicTocTrack®

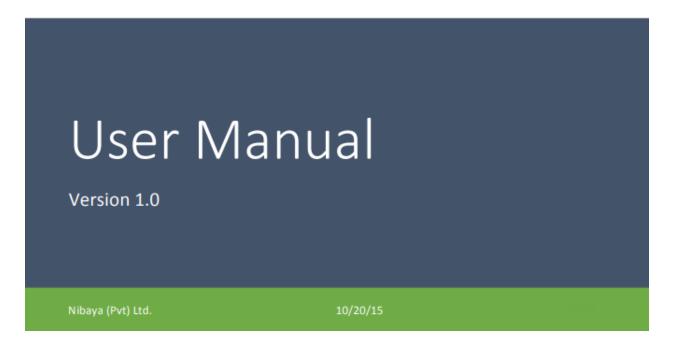
To download TicTocTrack®,

1. Follow the URL to get the download link for your desired platform

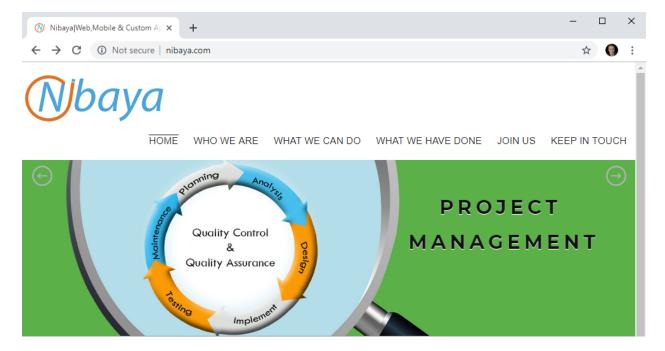
iOS – https://itunes.apple.com/us/app/tictoctrack/id1057481821
Android - https://play.google.com/store/apps/details?id=com.Nibaya. icTocTrack

Searching for that same "Nibaya" name on the TicTocTrack website turns up several different versions of the user manual:

www.tictoctrack.com.au

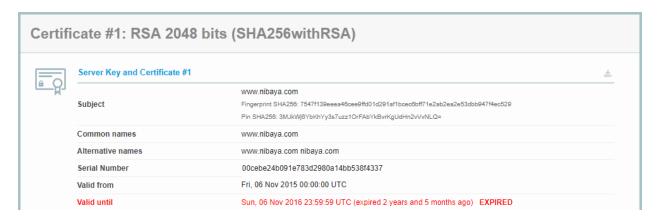


It turns out that <u>Nibaya</u> is a Sri Lankan software development company with a focus on quality control and quality assurance:

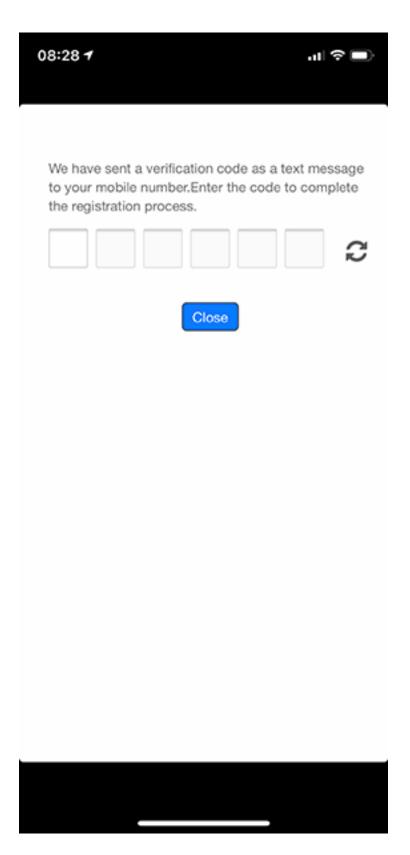


We're also told by the browser that they're "Not secure" which is not a great look

in this day and age. They do in fact have a certificate on the site, only thing is <u>it</u> <u>expired two and a half years ago</u> and they haven't bothered to renew it:



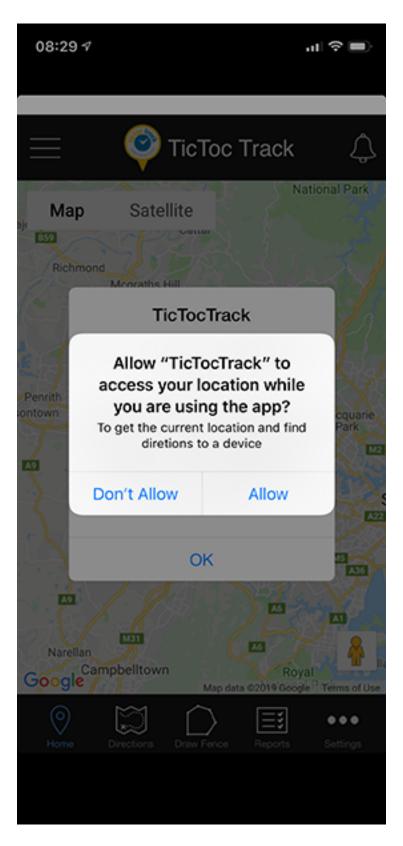
Moving on, there's a mobile phone number verification process which sends an SMS to my device:



Only thing is, the keyboard defaults back to purely alphabetical after every

character is typed so unless you pre-fill the field from the SMS (which iOS natively allows you to do), it's a bit painful. Again, it's all the little things.

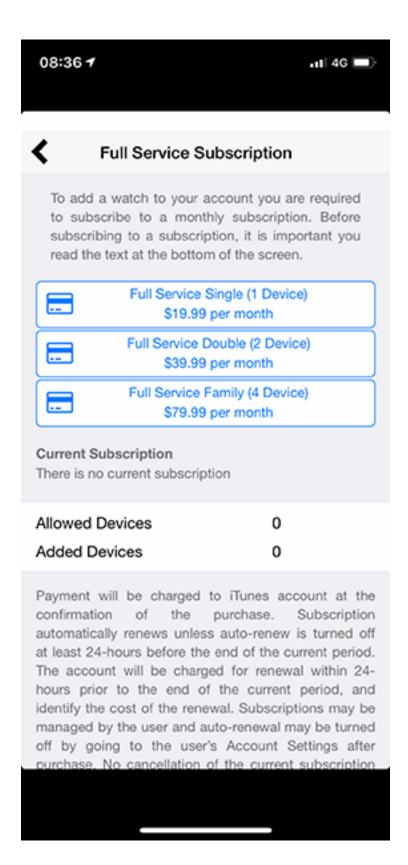
Following successful number verification, the app fires up and asks for access to location data:



Based on what I'd already read in the user manual, my location data can be used

to direct me to a child wearing the watch so requesting this seems fine for that feature to function correctly.

Next is the money side of things and we're looking at \$20 a month for the "Full Service Subscription":



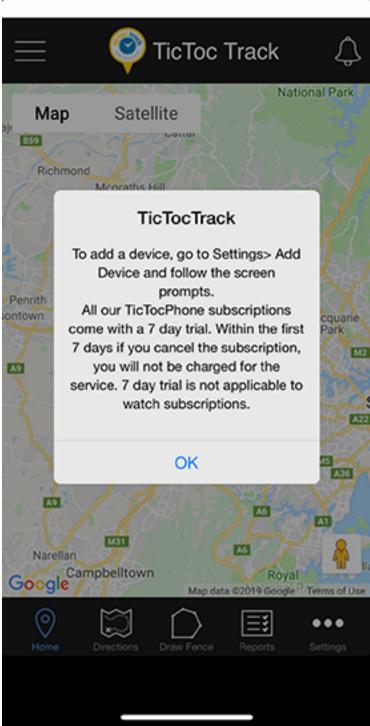
If I'm honest, I'm still a bit confused about what this entails. Is this for the

tracking service? Or for the Telstra SIM which it shipped with and is identically priced?



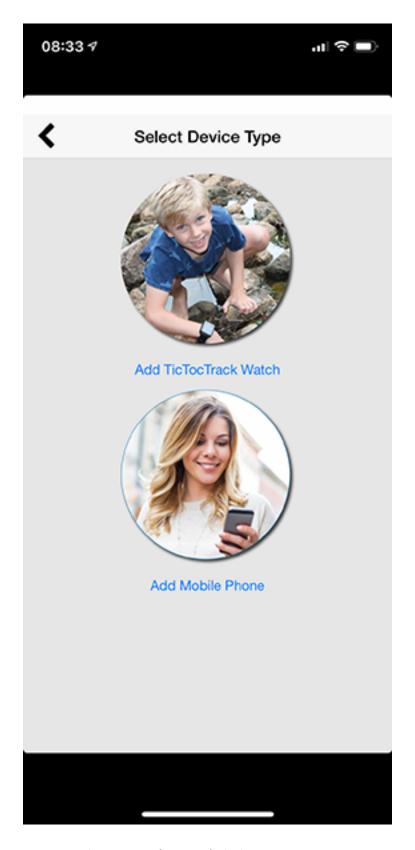
Or is it for both? I'm assuming both but then when I look at the <u>service plans</u> on the website, none of them are priced at \$19.99. Regardless, I take the \$20 option and move on:



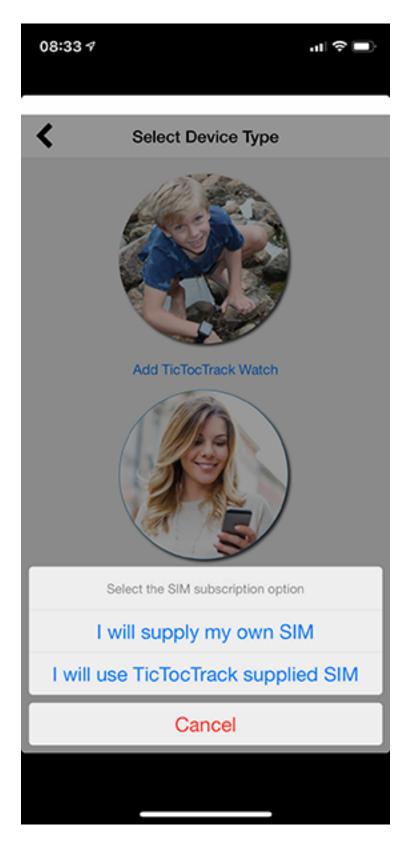


The adding a device bit I get - I'm going to need to pair the watch - but the

subscription bit further confuses me because I've literally just bought a subscription on the previous screen! For my purposes I don't see myself needing it for any more than 7 days anyway so I'm not too concerned, let's go and add that new device:

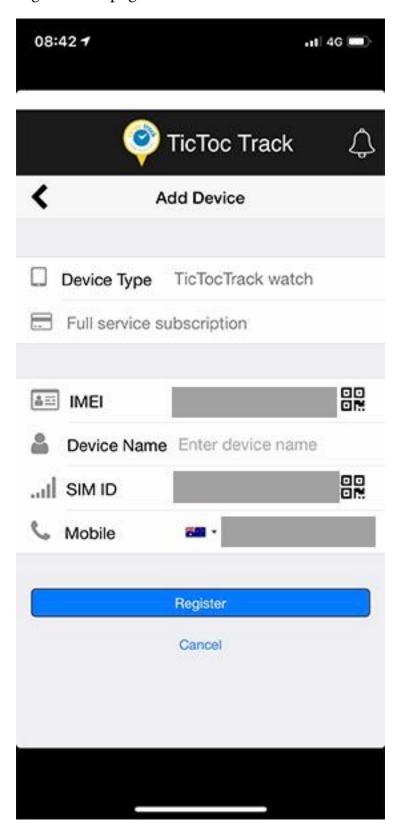


A new TicTocTrack watch it is:



And let's go with the supplied SIM which then leads us to the device and SIM

registration page:

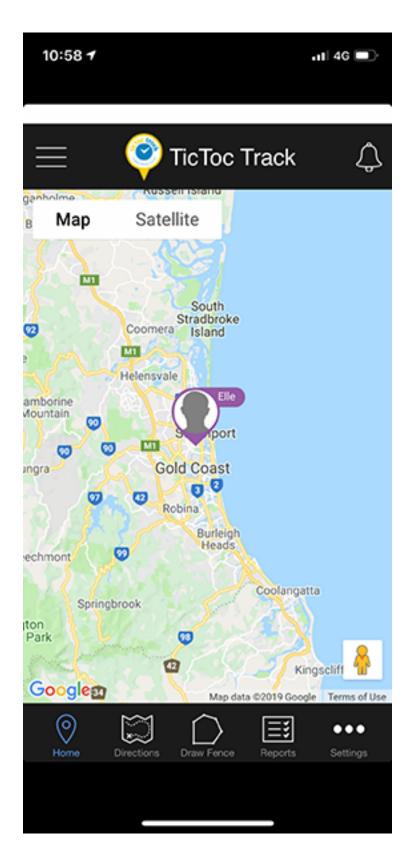


The IMEI is the identifier of the device itself (the watch) and that can be scanned off the barcode in the packaging. The SIM ID relates to the prepackaged SIM from Telstra, the barcode for which is under one of the grey obfuscation boxes in the earlier image. I call the device "Elle", register it and that's that.

Lastly, I insert the SIM into the watch (the metal flap for which opens in the opposite direction to the video tutorial and took me a good 5 minutes to work out for fear of breaking it), then drop it onto the power. Give it a couple of hours to charge, boot it up and shortly afterwards it's showing a 3G connection:

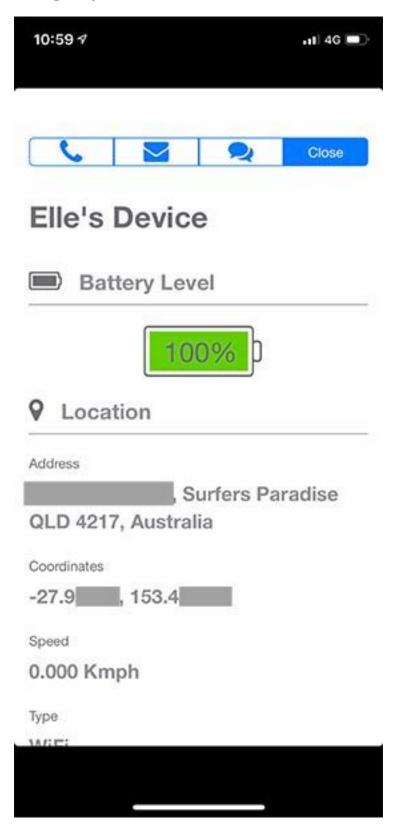


I give it a little time to sync to the TicTocTrack service then successfully find it in the app:



Drilling down on Elle's profile, I get an address and GPS coordinates which are

both pretty accurate:

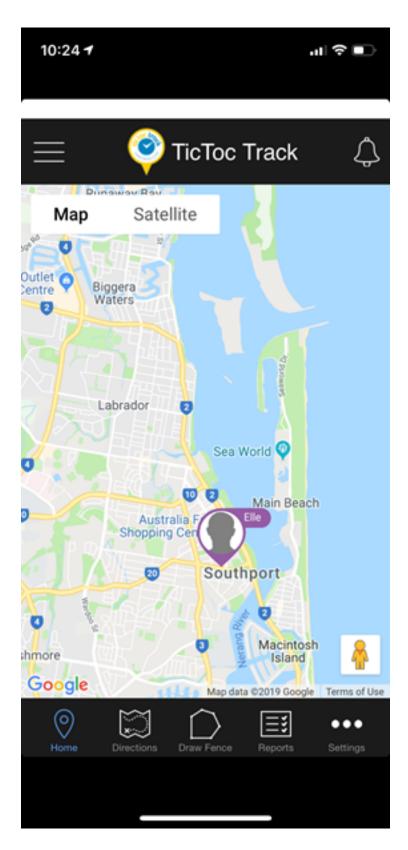


To its credit, the watch does a pretty good job of the setup and tracking process once you're past some of the earlier hurdles. At this stage, I now have a device which is broadcasting its location reliably and I can successfully see it in the app. I'm not going to go through other features such as the ability to send an SOS or make a call, at this stage all I really care about is that the watch is now tracking her movements.

The next day, we head off to tennis camp (it's school holiday time) with the TicTocTrack / Gator on her wrist:

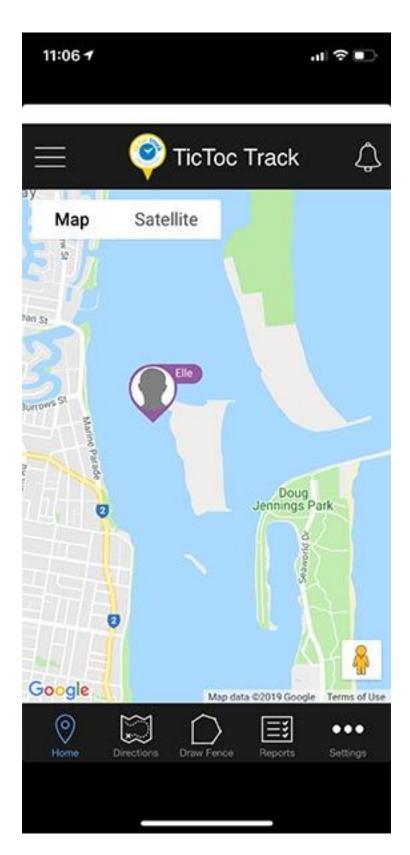


She isn't aware of why she has the watch, to her it's just a new cool thing she gets to wear. And it's pink so that's all boxes ticked. She's now at the local court whilst I (in my helicopter parent mode), am sitting at home watching her location on my device:



Safe in the knowledge that my little girl is in a place that I trust, I get back to

work. But someone else is also watching her location, someone on the other side of the world who is now able to track her every move - it's Ken. Not only is Ken watching, as far as TicTocTrack is concerned he's just taken her away:



She's no longer playing tennis, she's now in the water somewhere off Wavebreak

island. This isn't a GPS glitch; Ken has placed her four and a half kilometres away by exploiting an insecure direct object reference vulnerability in TicTocTrack's API. He's done this with my consent and only to my child, but you can see how this could easily be abused. It's not just the concept of making someone's child appear in a different location to what the parents expect, you could also have them appear exactly where the parents expect... when they're actually nowhere near there.

But these devices are about much more than just location tracking, they also enable 2-way voice communications just as you'd have on a more traditional cellular phone. This, in turn, introduces a far creepier risk - that unknown parties may be able to talk to your kids. In order to demonstrate this, I put the watch back on Elle and gave Pen Test Partners permission to contact her. Pay attention to how much interaction is required on her part in order for a stranger to begin talking to her simply by exploiting a vulnerability in the TicTocTrack service:



Even for me, that video is creepy. It required zero interaction because Vangelis was able to add himself as a parent and a parent can call the device and have it

automatically answer without interaction by the child. The watch actually says "Dad" next to a little image of a male avatar so a kid would think it was their father calling them:



This is precisely what the Germans were worried about when they banned the watches outright and when you watch that video, it seems like a pretty good move on their part.

The exploits go well beyond what I've already covered here too, for example:

wow I can also send SMS



so I can fake Dad and say "come meet me out of the tennis club"

Vangelis tix Stykas • Apr 13

and when I abduct her the sos button will call my phone...



facebook.com/travellingwith...

Vangelis tix Stykas • Apr 13



yet again marketing budget but no development or security budget

Vangelis tix Stykas • Apr 13

That link goes off to a Facebook post by an account called Travelling with Kids which very enthusiastically espouses the virtues of tracking them (it's not explicitly said, but the post appears to be promotional in nature):

The little wanderers were stoked to be going off to kids club at the Hard Rock Hotel Bali We have complete peace of mind knowing they're wearing their TicTocTrack watches, so they can call us at anytime and with GeoFencing we know their location

By now, I'm sure you can see the irony in the "peace of mind" statement.

The technical flaws go much further than this but rather than covering them here, have a read of the Pen Test Partners write-up which includes details of the IDOR vulnerability. Just to put it in layman's terms, here's the discussion I had with Vangelis about it:

Just so I understand correctly, initial auth is required with valid creds but then there's no validation that the family ID actually belongs to the auth'd user?







So you create your own user account then change the family ID which pulls someone else's data and gives you full control over their kid?







Being conscious that many people who don't normally travel in information security circles will read this, handling a vulnerability of this nature in a responsible fashion is enormously important. Obviously you want to remove the risk ASAP, but you also want to make sure that information about how to exploit it isn't made public beforehand. We religiously followed established best practices for responsible disclosure, here's the timeline with dates being local Aussie ones for me:

- 1. Saturday 6 April: Ken first contacts me about the watch. I order one that morning.
- 2. Tuesday 9 April: Watch arrives.
- 3. Wednesday 10 April: I set the account up.
- 4. Thursday 11 April: Elle wears the watch to tennis and we test "relocating" her.

5.Friday 12 April: Vangelis calls her and has the discussion in the video above. Ken privately discloses the vulnerability to TicTocTrack support that night.

6.Monday 15 April (today): TicTocTrack takes the service offline.

A couple of hours before publishing, I received a notification to the email address I signed up with as follows:



SERVICE INTERRUPTION TicTocTrack Platform

On Saturday 13 April, our customer service team received an email from Mr Ken Munro of Pen Test Partners - a limited liability partnership operating out of the UK. In this email, Ken alleged that there were security flaws in our TicTocTrack® software and that further detail on these alleged issues would be made public on Tuesday 16 April.

Although we have requested more information, no formal detail on these alleged flaws has been supplied to us by Ken Munro at this present time. Regardless, iStaySafe Pty Ltd (parent company of TicTocTrack®) takes claims like this seriously. To ensure our your data is kept safe and to further ensure the wellbeing of you and your loved ones, we will be taking the following action:

- We will be restricting user access to the TicTocTrack® application and service from 3.45pm today until such time as we can a) confirm the validity of the security flaws and b) fix the security flaws, if they exist
- iStaySafe will engage a trusted and accredited penetration testing provider to audit the TicTocTrack® security offering
- We will be offering refunds on subscriptions to our customers during the affected period, however, we look to have the issue resolved before this becomes necessary
- 4. We will respond to all of your enquiries as promptly as possible but please be patient with us if you experience delays during this period and rest assured we will come back to you as soon as we can.

We would like to confirm that to this day, there has **never** been a security breach that has lead to our customer's personal data being used for malicious purposes. Our dedicated team are constantly working to improve our software and make it as safe as possible for all our users.

As soon as a full technical assessment has taken place, conducted by a trusted, reputable and accredited penetration testing service, we will be releasing a transparent report which will detail what security issues (if any) were apparent, what steps we are taking and when - our aim is to have this issue resolved effectively and quickly so that you can resume your use of the

We hope to have the service back up and running by close of business Wednesday the 17th of April however will continue to keep you informed as we progress through our investigations.

We thank you for your continued loyalty and trust and can personally assure you that - the team at TicTocTrack and I Karen Cantwell, am doing everything we can to have the service back up and running as quickly as possible.

We apologise for any inconvenience this has caused. If you have any further questions, we are here to answer them - please email enquiries@tictoctrack.com.au or call 1300 872 256.

Kind regards,

Karen Cantwell,
Founder & CEO iStaySafe Pty Ltd

I'm in 2 minds about this message: on the one hand, they took the service down as fast as we could reasonably expect, being within a single business day so kudos to them on that. On the other hand, the messaging worries me in a number of ways:

Firstly, Ken didn't just "allege" that there were security flaws, he spelled it out. His precise wording was "The service fails to correctly verify that a user is authorised to access data, meaning that anyone can access any data, should they so wish". Anyone testing for a flaw of this nature would very quickly establish that changing a number in the request would hand over control of someone else's account thus proving the vulnerability beyond any shadow of a doubt. That word was used 3 times in the statement and it implies that they're unsubstantiated claims; they're clearly not. Which brings me to the next point:

Secondly, it wouldn't make sense to pull down the entire service if you weren't convinced there was a serious vulnerability. Many people allege there are security flaws in services but they don't generally go offline until they're proven. Clearly an incident like this has a bunch of downstream impact and acknowledging it publicly is not something you do on a whim. Either TicTocTrack was very confident in that accuracy of Ken's report (well beyond what "alleged" implies) or there were other factors I'm not aware of that drove them to rapidly pull the service.

Thirdly, the following statement was made without citing any evidence: "there has never been a security breach that has lead to our customer's personal data being used for malicious purposes". It's not uncommon to see a response like this following a security incident, but what it should read is "we don't know if there's ever been a security breach..." This vulnerability relied on an authenticated user with a legitimate account modifying a number in the request and the likelihood of that being logged in a fashion sufficient enough to establish it ever happened is extremely low. And if you were the kind of developers to log this sort of information, you'd also be the kind not to have the vulnerability in the first place!

Let's be perfectly clear - this is just one more incident in a series of similar ones impacting kids tracking watches and Gator in particular. What's infuriating about this situation is that not only do these egregiously obvious security flaws keep occurring, they're just not being taken seriously enough by the manufacturers and distributors when they do occur. There's no finer illustration of this than the statement Ken got when speaking to an agent over in his corner of the world:

UK agent for Gator said that they didn't have the money for security, as otherwise they couldn't afford a staff Xmas party

Is that really where we're at? Tossing up between exposing our kids in this fashion and beers at Christmas? If you're a parent ever considering buying one of these for your kid, just remember that quote. Inevitably, cost would have also been a major driver for TicTocTrack outsourcing their development to Sri Lanka, indeed it's something that Nabaya prides itself on:

We are Nibaya – teams that works with the customer in focus – Always.

"Always client focused" is our motto and we work with you for you, whether the task is big or small. Nibaya provides cost effective solutions to all your IT needs. Since we work with you, the client, in mind – you have nothing to fear to move forward with us.

I want to finish on a broader note than just TicTocTrack or Gator or even smart watches in general; a huge number of both the devices and services I see being marketed either directly at kids or at parents to monitor their kids are absolute garbage in terms of the effort invested in security and privacy. I mentioned CloudPets and VTech earlier on and I also mentioned spyware apps; by design, every one of these has access to data that most parents would consider very personal and, in many cases, (such as the photos older kids are often taking), very sensitive. These products are simply not designed with a security-orientated mindset and the development is often outsourced to cheap markets that build software on a shoestring. The sorts of flaws we're seeing perfectly illustrate that: CloudPets simply didn't have a password on their database and

both the VTech and TicTocTrack vulnerabilities were as easy as just incrementing a number in a web request. A bunch of the spyware breaches I referred to occurred because the developers literally published all the collected data to the internet for the world to see. How much testing do you think actually went on in these cases? Did nobody even just try adding 1 to a number in the request? Because that's all Ken needed to do; Ken can count therefore Ken can hack a device tracking children. Maybe I should give Elle a go at that, her counting is coming along quite nicely...

There's only one way I'd track my kids with GPS and cellular and that's with an Apple Watch. I don't mean to make that sound trivial either because we're talking about a \$549 outlay here which is a hell of a lot to spend on a kid's watch (plus you still need a companion iPhone), but Apple is the sort of organisation that not only puts privacy first, but makes sure they actually pay attention to their security posture too. As that Gator agent in the UK well knows, security costs money and if you want that as a consumer, you're going to need to pay for it.

I'll leave you with this thread I wrote up when first starting to look at the watch. It got a lot of traction and I'd like to encourage you to share it with your parenting friends on Twitter or via the one I also posted to Facebook.





I've been looking at a bunch of kid-related devices and services lately, mostly relating to how parents can monitor and control their activities. It's just consistently horrifyingly bad; FUD-ridden at best, massive privacy violations at worst (i.e. data accessible to the public).

3:28 PM · Apr 5, 2019

(i)





Replying to @troyhunt

The problem is that you've got a bunch of technically illiterate parents (understandable) being pushed things by schools that are influenced by marketers (much less understandable) and built with near zero focus on security (inexcusable).

3:28 PM · Apr 5, 2019







Replying to @troyhunt

You worried about your kids online? Talk to them. Browse the web with them. Introduce them to the wonders of the web on your terms and *physically* monitor them (you know, like exist together in the same room for a bit).

3:28 PM · Apr 5, 2019







Replying to @troyhunt

And accept that they're going to see porn. They're going to swear in chats. They're going to talk to people you don't like. And 90%+ of the time, they're more technically adept than their parents and will know how to hide it and circumvent the parental controls.

3:28 PM · Apr 5, 2019







Replying to @troyhunt

I'll talk to my kids all day long about this stuff, but I'll never install the sorts of software or buy the kinds of tracking devices I keep seeing peddled. These things are consistently absolute rubbish and they prey on scared and uninformed parents and teachers to get traction.

3:28 PM · Apr 5, 2019



Lastly, resources from the big OS vendors to help parents without putting kids even further at risk:

- Apple: https://apple.com/families/
- Microsoft: https://account.microsoft.com/family/about?ru=https:
 w2F%2Faccount.microsoft.com%2Ffamily...
- Google: https://families.google.com/familylink/

Comments

Hi Troy,

I enjoy reading your posts. They are long, but complete and really interesting. I never leave a comment because I honestly don't know what to add. I limit myself to share them with my friends.

But this time I need to add something. It is quite surprising to read this coming from you:

(...) but Apple is the sort of organisation that not only puts privacy first, but makes

sure they actually pay attention to their security posture too.

Really? Where is your blind and absolute confidence in Apple coming from? Ok, maybe there are spending a lot of money in security, and they do things "right" but it is not possible to know that for sure. Aren't you worry about Apple's backdoors? And more in general about any proprietary software when it comes to data safety?

Looking forward for the next post!

Troy: Yes, really. It's not blind confidence, it's based on a long track record they've maintained across a very broad range of products and services. It's also based on the heavy emphasis the organisation puts on privacy and you hear that message over and over again in their keynotes. Plus, you see it demonstrated in their design decisions and the implementations within their products. And as for back doors, there's *so* much scrutiny on these products that would be an enormously difficult thing to pull off undetected and it if it *was* detected, it would be enormously damaging for a company increasingly priding itself on privacy.

About back doors, it's not a questions if they have them or nor, it's a fact and they have admited them several times (I know it's old, but do you think Apple has changed since that?):

https://www.telegraph.co.uk...

And about encryption in the devices, it's true that since iOS 8 devices encrypt data with their own keys so that was a great move towards protecting consumer privacy and security. But what about iCloud? Apple encrypts iCloud data, but they encrypt it with their own key, not with your passcode key, which means that they are able to decrypt it to comply with government requests.

What I wanted to say with my previous comment, and also this one, is that is hard for me to trust Apple, their products and their OS, and that's in part because of reading your blog. Maybe I'm a bit paranoiac about this...

Troy: Geez, that's a leap! An 11 year old article about a feature clearly designed

to delete apps is a world away from a "backdoor"! Frankly, I don't have a problem with lawful access to iCloud, there's a valid discussion to be had about what the due process required to do that is, but that now becomes an issue around tech companies complying with local laws.

Time and again, you'll hear security pros pushing Apple over all others. Listen to Bob Lord talk about replacing DNC Androids with iPhones, for example. Or Runa Sandvik on protecting journos at the NYT with iPhones. And many others - there's a very clear trend there, and for good reason.

Thanks! I'll listen to Bob Lord and Runa Sandvik. Nevertheless, without any further reading/listening, I don't think it's a fair comparison Android vs iPhone... There are a lot of Androids out there, maybe even some old Android 2.4. To be fair, we should choose a specific phone from a specific company. And following Apple approach of doing everything (hw + sw) and selling the final product, only Google with their Pixel could compete in similar conditions I guess...

Anyway, I'll read about that. And thanks for taking your time to answer these comments. I can imagine how busy you are and still answering every single comment in your blog, kudos for that!

--

It's two entirely different issues. Having a big company like Apple collect personal info on your kid and then using it for marketing purposes is very different from a service that lets any old creep with IT-knowhow talk to your kid through their wristwatch, track their location in real time, AND make the same service actively misdirect the parents.

We should definitely be on guard about big corporations abusing personal info for less than than wholesome purposes, and what kind of access they have to our devices and data. But this is **not that**.

Troy: Except that Apple in particular is shaping themselves around being a privacy-centric company. Watch their keynotes - it's really interesting to see how

often they talk about things like usage data never leaving the device. I'd feel more comfortable with Apple providing services for me to locate my kids than any other organisation I can think of.

Allow me to clarify. We should be on guard about anyone collecting and misusing personal information. Small companies and big corporations alike. It wasn't meant as a stab against Apple in particular.

I merely wanted to point out that companies having access to your personal data (regardless of whether it's true or not), is not the same as a company exposing your kids location in real time to anyone with the inclination to look for it.

Epilogue

Geez, there's so much to unpack that happened after this blog post and it's a combination of security horror show and comedy gold. Let's start with their comms:

Companies responding to a security incident often put a lot of effort into spinning the story to lessen the impact on them. I get that, it's PR, but holy shit can it be cringeworthy. For example, amongst all the discussions that happened in the background of this incident, at one stage Ken received a statement from the CEO that came via a journalist. It said: "We ask the media to treat this incident for what it is, a couple of hackers trying to make a name for themselves". Ok, deep breath: Ken is both a lovely standout guy and a very accomplished professional with a heap of experience and public presence, including having previously been a TED speaker. I'm, well you can make up your own mind about me, but I'd also been to Congress so clearly had some credibility. Ken and I discussed just far under the bus we'd throw Karen for making such a stupid remark but because he's such a good bloke, he wanted to give her a chance to reconsider her position first. So, he reached out, politely put it to her that this wasn't a wise position to take and promptly learned that

the comments weren't hers, rather they were made by the PR company she'd engaged on her behalf. Now mind you, IMHO she'd still have to own those comments if they're speaking on her behalf, but I thought that was quite a fascinating outcome.

Another fun one was when a journo questioned her on the potential scale of the vulnerability: "How many Aussie families have been impacted by this?" A reasonable question very typical of journos, yet the response was that this information was confidential and wouldn't be shared. In the writeup that Pen Test Partners did about the incident, they included the vulnerable OData query to my record that looked like this:

Users?\$filter=(FamilyIdentifier eq 3497)

So, how many families were impacted? Hmmm... Moral of the story: if a number is meant to be a secret, don't put it in the bloody URL!

And then it got worse. 9 months after this story went public, their service was taken down, "fixed" and then restored, someone found that the original bug had returned. To demonstrate this, they sucked out a heap of personal data – including mine – and sent it over to me. I was in Oslo at the time doing the annual NDC things, so I didn't give it too much attention then, but immediately raised it with Karen. I ended up meeting her in person for a coffee in Brisbane during Feb 2020 and we discussed a whole heap of things related to her service in particular and the IoT landscape in general. She was nice, congenial and pretty much what I expected of a mum who starts up her own business selling kids tracking watching; well-intentioned but clearly well out of her depth on the security things.

TicTocTrack didn't disclose the 2020 breach. This bugged the hell out of me firstly because, well, it's a breach! Secondly, because it related to kids and there's a pretty broad consensus that anything to do with them deserves more care and attention than, say, your data from a forum devoted to cats. And finally, because my own data was in there. That's why I'm so confident they never disclosed the incident - because I never received a notification! In April

2020, several months after the incident, I wrote another blog about how reassuring words and good intentions don't mean good security which was a slightly roundabout way of pointing out the breach publicly.

The service is still up and running today, still talking about "a solution for peace of mind" and still leaning on the Aussie angles whilst not mentioning the Chinese or Sri Lankans. And I still wouldn't trust it, not one little bit.

PROJECT SVALBARD: THE FUTURE OF HAVE I BEEN PWNED

At the start of 2019, it felt like the world was closing in around me. As much as I loved working on HIBP, it had become a massive time sink and it was starting to rule my life. I was pouring huge amounts of time not just into maintaining the service but managing the data breach pipeline and then dealing with the fallout that would come after that. By "fallout" I mean the seemingly endless series of emails from journalists, people in the breach, security researchers, companies wanting to partner or pimp their wares, people wanting to kill me (there'd be the occasional death threat) and it was coming at me via every conceivable channel too. Emails, tweets, DMs, phone calls, encrypted chats; if someone could reach me by it, it was demanding my attention. The noise was deafening.

At the same time, my speaking career and publicly facing activities were also skyrocketing. I was travelling all over the world appearing at events, often as a featured speaker people had paid conference organisers good money to see. There'd be all this buzz leading up to the talk followed by an hour of entertainment (and I feel that's a reasonable characterisation; I wasn't doing tech talks, I was doing shows), then questions, selfies, handshakes and then... nothing. Silence. A lonely hotel room. I remember one particular occasion after Microsoft Ignite in Sydney a few months before writing this blog post where once I left the conference centre, I just felt like all the air had been sucked out of me. Perhaps it was after such a big build-up and so much effort going into the event that the aftermath was starker than it would have felt otherwise, but it was becoming a recurring theme.

I think this was my mid-life crisis. I was 42, married, kids, successful by any reasonable measure and healthy. Physically healthy, yet going through

emotional turmoil. The public veneer would have looked amazing and for the most part, it was, but privately, I was struggling. More than a year later I'd write about my divorce and with the benefit of hindsight, that situation influenced my decision to sell HIBP. There'd been a series of events that led to the relationship breaking down and like most breakups it's never usually one thing, it's many things over many years. The build-up of "things" had begun many years earlier and seriously accelerated through 2018, fuelled on by specific incidents. We decided together to pull the pin in April 2019, right at the same time I made the call to go through the HIBP M&A process.

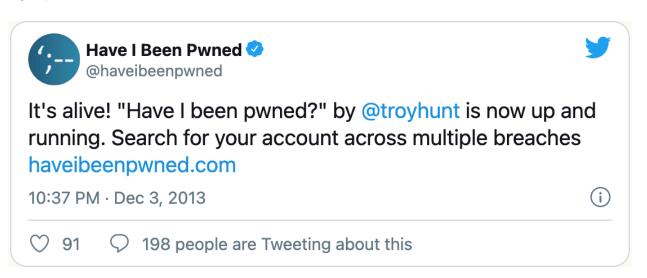
So now here I was, still in love with the project but struggling to find balance in my life.

11 JUNE 2019

Back in 2013, I was beginning to get the sense that data breaches were becoming a big thing. The prevalence of them seemed to be really ramping up as was the impact they were having on those of us that found ourselves in them, myself included. Increasingly, I was writing about what I thought was a pretty fascinating segment of the infosec industry; password reuse across Gawker and Twitter resulting in a breach of the former sending Acai berry spam via the latter. Sony Pictures passwords being, well, precisely the kind of terrible passwords we expect people to use but hey, actually seeing them for yourself is still shocking. And while I'm on Sony, the prevalence with which their users applied the same password to their Yahoo! accounts (59% of common email addresses had exactly the same password).

Around this time the Adobe data breach happened and that got me *really* interested in this segment of the industry, not least because I was in there. Twice. Most significantly though, it contained 153M other people which was a massive incident, even by today's standards. All of these things combined – the prevalence of breaches, the analysis I was doing and the scale of Adobe –

got me thinking: I wonder how many people know? Do they realise they were breached? Do they realise *how many times* they were breached? And perhaps most importantly, have they changed their password (yes, almost always singular) across the other services they use? And so Have I Been Pwned was born.



I'll save the history lesson for the years between then and today because there are presently 106 blog posts with the HIBP tag you can go and read if you're interested, let me just talk briefly about where the service is at today. It has almost 8B breached records, there are nearly 3M people subscribed to notifications, I've emailed those folks about a breach 7M times, there are 120k people monitoring domains they've done 230k searches for and I've emailed them another 1.1M times. There are 150k unique visitors to the site on a normal day, 10M on an abnormal day, another couple of million API hits to the breach API and then 10M a day to Pwned Passwords. Except even that number is getting smashed these days:





Pwned Passwords in ohaveibeenpwned is going from strength to strength - 16M requests in the last 24 hous with a cache hit ratio of 99.4% /cc olevapril



 \bigcirc 231 \bigcirc 45 people are Tweeting about this

Oh – and as I've written before, <u>commercial subscribers</u> that depend on HIBP to do everything from alert members of identity theft programs to enable infosec companies to provide services to their customers to protecting large online assets from credential stuffing attacks to preventing fraudulent financial transactions and on and on. And there are the governments around the world using it to protect their departments, the law enforcement agencies leveraging it for their investigations and all sorts of other use cases I never, ever saw coming (my legitimisation of HIBP post from last year has a heap of other examples). And to date, every line of code, every configuration and every breached record has been handled by me alone. There is no "HIBP team", there's one guy keeping the whole thing afloat.



@haveibeenpwned you guys are legends.

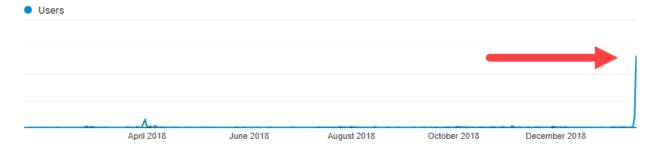
— CentristAgnostic (@BruvPeace) July 28, 2018

When I wanted an infographic to explain the architecture, I sat there and built the whole thing myself by hand. I manually sourced every single logo of a pwned company, cropping it, resizing it and optimising it. Each and every disclosure to an organisation that didn't even know their data was out there fell to me (and trust me, that's massively time-consuming and has proven to be the single biggest bottleneck to loading new data). Every media interview, every support

request and frankly, pretty much every single thing you could possibly conceive of was done by just one person in their spare time. This isn't just a workload issues either; I was becoming increasingly conscious of the fact that I was the single point of failure. And that needs to change.

It's Time to Grow Up

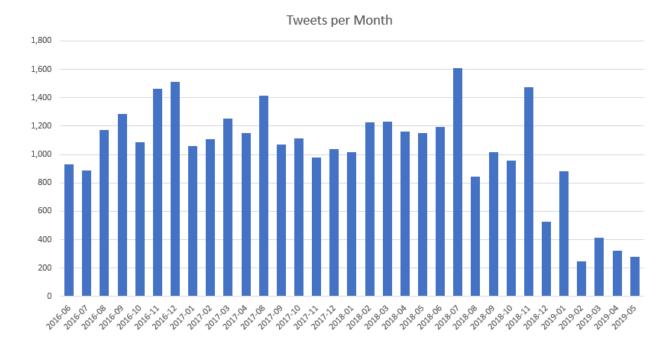
That was a long intro but I wanted to set the scene before I got to the point of this blog post: it's time for HIBP to grow up. It's time to go from that one guy doing what he can in his available time to a better-resourced and better-funded structure that's able to do way more than what I ever could on my own. To better understand why I'm writing this now, let me share an image from Google Analytics:



That graph is the 12 months to Jan 18 this year and the spike corresponds with the loading of the Collection #1 credential stuffing list. It also corresponds with the day I headed off to Europe for a couple of weeks of "business as usual" conferences, preceded by several days of hanging out with my 9-year old son and good friends in a log cabin in the Norwegian snow. I was being simultaneously bombarded by an unprecedented level of emails, tweets, phone calls and every other imaginable channel due to the huge attention HIBP was getting around the world, and also turning things off, sitting by a little fireplace in the snow and enjoying good drinks and good conversation. At that moment, I realised I was getting very close to burn-out. I was pretty confident I wasn't actually burned out yet, but I also became aware I could see that point in the not too distant

future if I didn't make some important changes in my life. (I'd love to talk more about that in the future as there are some pretty significant lessons in there, but for now, I just want to set the context as to the timing and talk about what happens next.) All of this was going on at the same time as me travelling the world, speaking at events, running workshops and doing a gazillion other things just to keep life ticking along.

To be completely honest, it's been an enormously stressful year dealing with it all. The extra attention HIBP started getting in Jan never returned to 2018 levels, it just kept growing and growing. I made various changes to adjust to the workload, perhaps one of the most publicly obvious being a massive decline in engagement over social media, especially Twitter:



Up until (and including) December last year in that graph, I was tweeting an average of 1,141 times per month (for some reason, Twitter's export feature didn't include May and June 2017 and only half of July so I've dropped those months from the graph). From Feb to May this year, that number has dropped to 315 so I've backed off social to the tune of 72% since January. That may seem like a frivolous fact to focus on, but it's a quantifiable number that's directly attributable to the impact the growth of HIBP was having on my life. Same again

if you look at my blog post cadence; I've religiously maintained my weekly update videos but have had to cut way back on all the other technical posts I've otherwise so loved writing over the last decade.

After I got home from that trip, I started having some casual conversations with a couple of organisations I thought might be interested in acquiring HIBP. These were chats with people I already knew in places I respected so it was a low-friction "put out the feelers" sort of situation. It's not the first time I'd had discussions like this – I'd done this several times before in response to organisations reaching out and asking what my appetite for acquisition was like – but it was the first time since the overhead of managing the service had gone off the charts. There was genuine enthusiasm which is great, but I quickly realised that when it comes to discussions of this nature, I was in well over my head. Sure, I can handle billions of breached records and single-handedly run a massive online data breach services that's been used by hundreds of millions of people, but this was a whole different ballgame. It was time to get help.

Project Svalbard

Back in April during a regular catchup with the folks at KPMG about some otherwise mundane financial stuff (I've met with advisers regularly as my own financial state became more complex), they suggested I have a chat with their Mergers and Acquisition (M&A) practice about finding a new home for HIBP. I was comfy doing that; we have a long relationship and they understand not just HIBP, but the broader spectrum of the cyber things I do day to day. It wasn't a hard decision to make - I needed help and they had the right experience and the right expertise.

In meeting with the M&A folks, it quickly became apparent how much support I really needed. The most significant thing that comes to mind is that I'd never really taken the time just to step back and look at what HIBP actually does. That

might sound odd, but as it's grown organically over the years and I've built it out in response to a combination of what I think it should do and where the demand is, I've not taken the time to step back and look at the whole thing holistically. Nor have I taken enough time to look at what it could do; I'm going to talk more about that later in this post, but there's so much potential to do so much more and I really needed the support of people that specialise in finding the value in a business to help me see that.

One of the first tasks was to come up with a project name for the acquisition because apparently, that's what you do with these things. There were many horribly kitschy options and many others that leaned on overused infosec buzzwords, and then I had a thought: what's that massive repository of seeds up in the Arctic Circle? I'd seen references to it before and the idea of a huge vault stockpiling something valuable for the betterment of humanity started to really resonate. Turns out the place is called Svalbard and it looks like this:



Also turns out the place is part of Norway and all these things combined started

to make it sound like a befitting name, beginning with the obvious analogy of storing a massive quantity of "units". There's a neat video from a few years ago which talks about the capacity being about a billion seeds; not quite as many records as are in HIBP, but you get the idea. Then there's the name: it's a bit weird and hard to pronounce for those not familiar with it (although this video helps), kinda like... pwned. And finally, Norway has a lot of significance for me being the first international talk I did almost 5 years ago to the day. I spoke in front of an overflowing room and as the audience exited, every single one of them dropped a green rating card into the box.





The feedback after @troyhunt's talk #ndcoslo



3:08 AM · Jun 6, 2014

<u>(i)</u>

That was an absolute turning point in my career. It was also in Norway this January that HIBP went nuts as you saw in the earlier graph. It was there in that little log cabin in the snow that I realised it was time for HIBP to grow up. And

by pure coincidence, I'm posting this today from Norway, back again for my 6th year in a row of NDC Oslo. So as you can see, Svalbard feels like a fitting name?

My Commitments for the Future of HIBP

So what does it mean if HIBP is acquired by another company? In all honesty, I don't know precisely what that will look like so let me just candidly share my thoughts on it as they stand today and there are a few really important points I want to emphasise:

- 1.Freely available consumer searches should remain freely available. The service became this successful because I made sure there were no barriers in the way for people searching their data and I absolutely, positively want that to remain the status quo. That's number 1 on the list here for a reason.
- 2.**I'll remain a part of HIBP.** I fully intend to be part of the acquisition, that is some company gets me along with the project. HIBP's brand is intrinsically tied to mine and at present, it needs me to go along with it.
- 3.**I want to build out much, much more capabilities wise.** There's a heap of things I want to do with HIBP which I simply couldn't do on my own. This is a project with enormous potential beyond what it's already achieved and I want to be the guy driving that forward.
- 4.**I want to reach a much larger audience than I do at present.** The numbers are massive as they are, but it's still only a tiny slice of the online community that's learning of their exposure in data breaches.
- 5. There's much more that can be done to change consumer behaviour. Credential stuffing, for example, is a *massive* problem right

now and it only exists due to password reuse. I want HIBP to play a much bigger role in changing the behaviour of how people manage their online accounts.

- 6.**Organisations can benefit much more from HIBP.** Following on from the previous point, the services people are using can do a much better job of protecting their customers from this form of attack and data from HIBP can (and for some organisations, already does) play a significant role in that.
- 7. There should be more disclosure and more data. I mentioned earlier how responsible disclosure was massively burdensome and Svalbard gives me the chance to fix that. There's a whole heap of organisations out there that don't know they've been breached simply because I haven't had the bandwidth to deal with it all.

In considering which organisations are best positioned to help me achieve this, there's a solid selection that are at the front of my mind. There's also a bunch that I have enormous respect for but are less well-equipped to help me achieve this. As the process plays out, I'll be working with KPMG to more clearly identify which organisations fit into the first category. As I'm sure you can imagine, there are some very serious discussions to be had: where HIBP would fit into the organisation, how they'd help me achieve those bullet-pointed objectives above and frankly, whether it's the right place for such a valuable service to go. There are also some major personal considerations for me including who I'd feel comfortable working with, the impact on travel and family and, of course, the financial side of the whole thing. I'll be honest - it's equal parts daunting and exciting.

Last week I began contacting each stakeholder that would have an interest in the outcome of Project Svalbard before making it public in this blog post. I explained the drivers behind it and the intention for this exercise to make HIBP not just

more sustainable, but also for it to make a much bigger impact on the data breach landscape. This has already led to some really productive discussions with organisations that could help HIBP make a much more positive impact on the industry. There's been a lot of enthusiasm and support for this process which is reassuring.

One question I expect I'll get is "why don't I turn it into a more formal, commercially-centric structure and just hire people?" I've certainly had that opportunity for some time either by funding it myself or via the various VCs that have come knocking over the years. The main reason I decided not to go down that path is that it *massively* increases my responsibilities at a time where I really need to reduce the burden on me. As of today, I can't just switch off for a week and frankly, if I tried even for a day I'd be worried about missing something important. In time, building up a company myself might allow me to do that but only after investing a substantial amount of time (and money) which is just not something I want to do at this point.

Summary

I'm enormously excited about the potential of Project Svalbard. In those early discussions with other organisations, I'm already starting to see a pattern emerge around better managing the entire data breach ecosystem. Imagine a future where I'm able to source and process much more data, proactively reach out to impacted organisations, guide them through the process of handling the incident, ensure impacted individuals like you and me better understand our exposure (and what to do about it) and ultimately, reduce the impact of data breaches on organisations and consumers alike. And it goes much further than that too because there's a lot more that can be done post-breach, especially to tackle attacks such as the huge rate of credential stuffing we're seeing these days. I'm really happy with what HIBP has been able to do to date, but I've only scratched the surface of potential with it so far.

I've made this decision at a time where I have complete control of the process. I'm not under any duress (not beyond the high workload, that is) and I've got time to let the acquisition search play out organically and allow it to find the best possible match for the project. And as I've always done with HIBP, I'm proceeding with complete transparency by detailing that process here. I'm *really* conscious of the trust that people have put in me with this service and every single day I'm reminded of the responsibility that brings with it.





@troyhunt I just wanted to say I think you're doing God's work with @haveibeenpwned. I've used it with every company I've worked for so far.

9:06 AM · Jun 2, 2019



HIBP may only be less than 6 years old, but it's the culmination of a life's work. I still have these vivid memories stretching back to the mid-90's when I first started building software for the web and had a dream of creating something big; "Isn't it amazing that I can sit here at home and write code that could have a real impact on the world one day". I had a few false starts along the way and it took a combination of data breaches, cloud and an independent career that allowed me the opportunity to make HIBP what it is today, but it's finally what I'd always hoped I'd be able to do. Project Svalbard is the realisation of that dream and I'm enormously excited about the opportunities that will come as a result.

Comments

Thank you so much, Troy for everything you have done so far - especially with HIBP, it is

such a valuable resource and I am really glad that you will stay involved with it. Good luck with this and I hope it all goes to plan and continues to run, however it does so. Also hope it gets you a bit more downtime! Good luck and thank you! :-)

I'm not entirely comfortable with something as important as HIBP belonging to a for-profit company, to be honest (having a bus factor of 1 isn't good either, as is the point of your post).

Have you approached the EFF or some similar nonprofit? Another option I can think of is to ensure HIBP ends up in an organization like a public-benefit corporation, where their stated goals are aligned with yours.

Came down to the comments chiefly to post this.

Once someone else owns it, Troy gets no actual say in the future of it, no matter what he wants for it (no matter what the acquiring corporation says -- we have plenty of examples of that). That's very dangerous for a resource like HIBP.

If it is acquired, I necessarily have to completely discard every scrap of trust and faith I've put in the service and start from 0.

There are also going to be an enormous number of questions that most companies won't want to answer, like how secure their systems are, how secure these data breaches will be kept internally, and who exactly will have access to the data.

Troy can easily police himself, but it is actually impossible to police an entire org (HIBP's necessary existence itself proves that).

Troy: I don't agree with all those comments (such as my having no say in the future of HIBP post-acquisition), but I will say this: I get *complete* say over who acquires it in the first place and a massive part of my decision-making process will be to choose the organisation that best aligns with my view of how

this service should be run. That's not just a commitment to you and everyone else that has come to trust HIBP, it's a commitment to myself because I want to be around other people that share the same view of how data breaches should be handle. I cannot overstate how important that is to me.

I'm definitely not saying that you have any intent otherwise. I fully believe that you are fully committed to the future of HIBP and ensuring the ethical behaviour of the company that acquires it.

But I *AM* saying that we've had numerous past examples of founders who were forced out or shifted in the organisation so that they lose control over their product. Even when companies we trust (ed) were involved.

It is so endemic to the process that Harvard Business Review wrote an article: <a href="https://https:

You also won't have control over the new corporation's direction, so if the corporation you sell to gets sold off to, say, Palantir...

I'm sure there may be ways to avoid this, and I'm sure that if there are, you'll find them. You've absolutely earned our trust and our faith.

I just don't believe that trust and faith are transferable.

Troy: I don't disagree with most of what you've said in this comment; it's possible a company may change direction in the future. All I can do at this stage is make the best possible choice with the information I have and do everything I can to ensure my vision for the service is fulfilled.

Hey Troy it appears to be a rollercoaster ride for you, love your blogs and advice you give,

and the way you explain everything, it's what keeps me in the infosec business, if i could articulate as well as you can and make it accessible to worker bees and C-Suite execs as well I would be a happy man. You should be very proud of what you have accomplished with HIBP and most people who comment I hope appreciate the effort you have gone to to explain your reasoning and thought process the way you always do. The very best of luck with this project and if financially you do well - brilliant it's your business and yours alone.

I loved your post on guidance for developers or infosec people looking to go independent and what this takes and entails, you are an inspiration and would love to be in a position to follow in your footsteps, alas the conditions are not right for me.

I can empathise with your pfizer days as I currently work as an ISO for an ecommerce/moto company dealing day to day with PCI, GDPR and just starting to touch on ISO27001, and doing it without any formal qualifications, all i have is 9 years experience to my name. Prior to this I did a couple of years as a Perl developer and a web designer before that so not really anything substantial such as yourself to base an independent career on.

My work is 12 miles away from home and because medically I cannot drive I have to suffer 2+ hour commutes each way 5 days a week so i get the 9-5 life is depressing what makes it worse is that while i am so passionate about infosec, the decision makers don't really get it so it's a case of chipping away a little bit at a time.

Anyway enough about poor poor me, really thrilled and looking forward to watching this journey unfold - good for you decent blokes come good, there is hope for us all.

Epilogue

I'd moved out of home only a week before writing this. It was my dream home in a location I'd idolised since I was a kid, something I'd finally been able to achieve after so many years of hard work and it was a place I'd never wanted to leave. I'd headed back to London and onto the big stage to receive my Hall of Fame award from Infosecurity Europe and again, it felt like the best of times

and the worst of times. I'd spend the next 11 weeks hopping between hotels and Airbnbs in Europe, the Middle East and the US, simultaneously going through the most stressful business process I could imagine whilst dealing with the most stressful personal circumstances I'd ever experienced. Plenty of other people have gone through the latter, including plenty of people who'll read this and relate. The former, however, was exceptional so let me share a bit of what it looked like behind the scenes:

The demands on my time whilst going through the M&A process were enormous. Initially it was bringing the KPMG guys up to speed on what on earth this thing was that I'd created, then it was trying to take that and turn it into something we could market. What we quickly realised is that it wasn't HIBP we were selling, rather it was me and HIBP would come along for the ride. Nobody wanted HIBP, not on its own at least, they wanted me as the public face of it. Later on, that would become a cornerstone of every single offer that was put forward and I was going to be locked into the acquiring organisation for years to come. So, we were simultaneously building up a pitch for HIBP and for me personally, both as assets to be traded which started to make me feel very, very uncomfortable.

But it was the questions from bidders which really stung time wise. There was a whole bunch of very reasonable stuff around the data breach pipeline, technical details and financial info, but there was also a huge amount from lawyers and M&A people which was just... mind numbing. Imagine a big tech company (and whichever one you're thinking of, there's a pretty solid chance they were one of the ones bidding) and think about how many acquisitions of all shapes and sizes they make every year. This is cookie-cutter stuff that they repeat over and over again so they have a process. And forms. And standard questions to be asked and I got hit with all of it. I couldn't delegate any of it because it was only me and all the time I spent doing that, I wasn't doing other things. It sucked, and it defined my life from April 2019 until March the following year when finally, it ended.

PROJECT SVALBARD, HAVE I BEEN PWNED AND ITS ONGOING INDEPENDENCE

I wrote this post at a time where life was turning a corner. I'd gotten back to Australia the month before and moved back into my own home (my ex had decided to move out). Now in a relationship with Charlotte, she'd moved in with me and we were establishing our own new "normal". The HIBP deal had just fallen over which was a shock but equally, a relief, like the pain was finally over. I was over caring about the money in the deal and frankly, it was seriously complicating a number of things on the personal front. All I wanted to do was to settle down and get some consistency in life.

I was still suffering from the pace of 2019 when I wrote this post. Relationship breakdown. Business sale. Huge amounts of travel. I was burned out and on reflection, that tone comes through in the blog post if you look beyond the words themselves. I wasn't ready to talk about a lot of that at the time (the next post in this book 4 months later on sustaining performance covers that) but it significantly influenced my words and my writing style. There's a steely façade in this post, but it's sitting on top of a lot of pain.

03 MARCH 2020

his is going to be a lengthy blog post so let me use this opening paragraph as a summary of where <u>Project Svalbard is at</u>: Have I Been Pwned is no longer being sold and I will continue running it independently. After 11 months of a very intensive process culminating in many months of exclusivity with a party I believed would ultimately be the purchaser of the service, unexpected changes to their business model made the deal

infeasible. It wasn't something I could have seen coming nor was it anything to do with HIBP itself, but it introduced a range of new and insurmountable barriers. So that's the tl;dr, let me now share as much as I can about what's been happening since April 2019 and how the service will operate in the future.

In the Beginning, There Were 141 Companies

According to the lock screen, I took the photo below at 04:49 on the 24th of July last year. I was in yet another bland, nondescript hotel room, drinking bad coffee in an attempt to stave off the jet lag. I'd arrived in San Francisco a few days earlier after barely making my connection in Helsinki, literally running through the airport. My bag hadn't made it. I was tired, alone, emotional and if I'm honest, at an all-time low. I snapped this pic to remind me how much energy I was pouring into the project when I came out the other side, whatever the outcome may be.



One day I'd really like to turn this whole experience into a conference talk because it's a fascinating story, but for now I want to try and give a sense of just how intense the last 11 months has been, starting with the heading above. Per the Project Svalbard announcement blog post, I engaged KPMG to run the merger and acquisition (M&A) process for me. Between their outreach to suitable organisations and the inbound requests from others after writing the announcement blog post, we spoke to a total of 141 different companies from around the globe. (The total number of organisations under consideration was actually *significantly* higher than that, but we culled all those we didn't consider "Tier 1" or in other words, highly likely to be a good fit for HIBP.) These were companies spanning all sorts of different industries; big tech, general infosec, antivirus, hosting, finance, e-commerce, cyber insurance - I could go on. The point is the net was cast very wide.

We whittled the original 141 companies down to the 43 that were best aligned to the goals I outlined in the original blog post. As I've said throughout this process, the decision around who I wanted to entertain as a bidder for the service was always going to be mine and mine alone so I culled companies that I didn't believe should have responsibility for the sort of data HIBP has, that wouldn't shepherd the service in the direction I believed it should go or were simply companies that I didn't want to work for. That last point is critical - it was repeated over and over again by every single organisation we discussed it with that a sale of HIBP was also a sale of *me* for many years to come. I would be an employee. I'd fly the company flag. I'd need to support their vision. That was only ever going to happen with a company that I wanted to devote many years of my life to. So, in late July I flew to San Francisco and spent a couple of weeks meeting with those 43 companies, KPMG guys in tow. It all felt a bit, well...



Seriously, I vividly remember after one early meeting on-premise with a tech company, walking out of the building with the KPMG guys laughing about how much it felt like an episode of Silicon Valley! Over and over again, we'd go to

these meetings and sit across the table from characters that could have come straight out of the show. The Russ Hannemans, the Gavin Belsons, the Lori Breens and here's me, feeling all Richard Hendricks. I hope I was a bit more articulate than Richard, but I was someone fronting up and presenting my pride and joy to strangers who I hoped would share the same enthusiasm for it that I did. I make the Silicon Valley comparison only partly tongue-in-cheek because it was absolutely uncanny how true the experiences tracked to the comedy.

The 43 companies we met with all received an <u>information memorandum</u> (IM):

An Information Memorandum (IM) is a package of documents created by business owners for prospective buyers. The primary mandate of an Information Memorandum is to motivate potential investment into your business. Although this package is designed to draw the interest of prospective buyers, it dually serves the purpose of transparency. Owners should avoid exaggeration, and aspire to disclose any information that will materially affect the value of the company.

In other words, you're trying to provide as much information as possible about the business so that potential purchasers can simultaneously understand what it does, see where the future potential is, foresee any risks, value it and ultimately put forward a bid.

Tangentially to the IM, one thing that worked in my favour when it came to providing information about how HIBP operates is that because I've run it with such transparency for so long, a lot of questions had already been answered publicly. For example, I was regularly asked if I'd ever received any legal threats which is apparently pretty normal for any M&A process, but you can imagine why it'd be particularly interesting when dealing with a heap of data originally obtained via illegal methods. No, I've never received any legal threats of substance (for example, I've never received a letter from a legal representative threatening to take action against me), the closest example I could think of I'd already included in one of my talks (deep-linked to the "high priced lawyers" point in the video):



I never heard anything further from Adult Fan Fiction?

The IM description goes on:

Information Memorandums tend to be very exhaustive as they should include items relating to the financial standing, assets and liabilities, business description, market position, clients, strategies and promotion methods, markets served, etc. of the company.

"Exhaustive" doesn't even begin to explain the effort that went into the Project Svalbard IM. We spent *months* preparing the document, regularly working until all hours to flesh it out as comprehensively as possible.



Looking back through the IM now, it had everything from traffic stats to revenue to assets to debts (none!) to customers to noteworthy events since conception to a slide on "Industry Tailwinds" talking about how big cyber is becoming (that hurt a little bit to put my name on, so much cyber...). Here's a page from it that was intended to pimp my own personal credentials:

Founder profile - Troy Hunt



HIBP's founder, Troy Hunt, is a renowned information security subject matter expert with over 20 years' experience in software development



In 2017, Troy appeared as an expert before Congress in the United States in a hearing on Identity Verification in a Post-Breach World

This experience and endorsement as an expert carries significant intangible value to the HIBP service and the founder behind it



"Troy is recognised for his contribution to the advancement of information security good practice as an industry advocate through his security research, public education and outreach work and for creating HIBP, enabling non-technical users to discover whether their data has been compromised"

Background and experience

- Troy often appears as a keynote speaker at software and security events around the world
- His personal blog has a considerable following and features posts on experiences with (and learnings from) various information security products and developments
- In 2017, he was requested to appear before Congress in the United States for a hearing on Identity Verification in a Post-Breach World
- Troy currently holds Microsoft's Most Valuable Professional and Regional Director titles as representations of his standing and influence in the community (Troy is not employed by Microsoft)
- He received the 2018 AusCERT Individual Excellence in Information Security Award and also received the Grand Prix Prize at the European Security Blogger Awards in the same year
- In 2019, Troy was inducted into the InfoSecurity Hall of Fame
- He has produced a number of training materials focused on information security and he regularly facilitates workshops designed to introduce software developers to secure coding practices
- Troy's reputation in the community and relationships have a unique and
 positive impact on HIBP through his ability to gather data from an array of
 reliable sources (e.g. breached organisations, white-hat contributors)



©2019 KPMG Corporate Finance LLC, a Delaware limited liability company. Member FINRA and SIPC, KPMG Corporate Finance LLC is a subsidiary of KPMG LLP, a Delaware limited liability company. Member FINRA part of the LLP, and the LLS member Firm of the KPMG network of indexendent member firms of the Microscopic Control of the Control of

ocument Classification: KPMG Confidential

This was another really unexpected part of the experience - how people perceived me personally and put a value on my brand. I was really conscious that the companies weren't bidding for HIBP, *they were bidding for me running HIBP* so a significant part of the purchase price was quite literally a dollar figure on my head. A little while back I had a discussion with someone who wanted to collaborate as they weren't getting the traction they wanted when pitching their own product to major tech firms. He made the following comment about trust:

I kid you not, was in a meeting at [big tech company] HQ in [HQ location] and a comment was made to the effect that "there is only one service they trust as a white hat (Troy and HIBP) and I'm like "fuck how does one guy corner the market on trust?"

This is what the organisations bidding on HIBP were buying: trust in me. Anyone can cobble together a website with some APIs and load in a ton of data breaches, but establishing *trust* is a whole different story. Trust in the way I run

the service is an absolutely pivotal part of HIBP and it's something I built organically rather than setting out to earn it, now here I was with big companies putting a value on it. That felt weird in a way I've never experienced before, certainly not like in times gone by where I'd interviewed for jobs. But then it was also an exciting time where I'd walk into a meeting with a company and they'd be so enthusiastic to meet me in person after following me for years so we'd do selfies, hand out HIBP stickers and then settle into serious business discussions. It was surreal.

Reflecting on the Process in August

Time and time again throughout Project Svalbard, I questioned whether it was the right thing to do. The *motives* were right in that it was first and foremost for the sustainability of the project so I wasn't concerned about that, but was selling HIBP genuinely the best path forward? Was this the future I wanted? Of course, I'd considered all that before making the decision to go down this path, but nothing could prepare me for the actual emotions felt once I was eyeball-deep in the M&A process.

I had a seminal moment just after all the San Francisco meetings as I was making my way over to the Black Hat and Defcon conferences in Vegas. The inperson meetings had wrapped up but that didn't stop a never-ending stream of teleconferences (I destroyed the batteries on a set of AirPods just from Project Svalbard teleconferences). I was parked in my rental car talking to the guy who'd be my boss if the large tech company he worked for emerged as the successful bidder. He asked a question - a perfectly reasonable interview question - but it sent chills down my spine:

So Troy, explain to me what your perfect day in the office would look like.

I kid you not, the immediate thought that popped into my mind was "I get up, get on my jet ski then do whatever the fuck I want". I can't remember exactly

how I answered the question, but I *can* remember how it made me feel and it was pretty damn uncomfortable. The last "job" I had I absolutely hated by the end of it. My boss was an arsehole (there was broad consensus on that noun), but I stuck it out and dealt with it until circumstances were such that there was a better path forward; ultimately, a redundancy with a nice payout (I cover this in my Hack Your Career talk). I *love* my life of independence and whilst I was prepared to work for a company again, it had to be *the right company* and this just felt... wrong. Many of them felt wrong.

Only a day later I received an email that reminded me how important HIBP was not just for me, but for an untold number of other people:

Sincerest of gratitude



2 Aug

Good afternoon Troy,

My name is Cody, and I am a System Administrator for a non-profit Cancer Research company. I'm extremely passionate about ITSec and currently working my way to break into the Red Teaming industry. I heard you on Darknet Diaries, and I just wanted to express my complete gratitude for https://haveibeenpwned.com/. I have personally relied on your website for years regarding my personal info for data breaches. Until I heard you speak, I never even considered just entirely how much work went into upkeep and maintaining the website. I really just wanted to say that I appreciate the effort you have put into the website.

I asked Cody at the time if he'd mind me sharing this at a later date then dropped it into this draft blog post. I didn't know what the post would say at the time, it was either going to announce a successful bidder or announce that HIBP would remain an independent project. Either way, this email was going in there to reinforce how important the trust of those who use HIBP is to me. Whatever the outcome, I wasn't going to do anything to let the Codys of the world down.

Then I got to Vegas. I wandered the conference halls with <u>Scott Helme</u> for a week and time again had complete strangers come up to me and thank me for HIBP (they also constantly asked <u>who the fuck the other guy was</u> which brought a much-needed smile to my face?). It happened dozens of times, often with much excitement, selfies and <u>exchanges of radio waves across Defcon badges</u>. After one such encounter, I added the following to the draft blog post with Cody's email and I'm reproducing it here precisely as I wrote it in the midst of the M&A process 7 months ago now:

I remember one discussion in particular where the guy was talking so sincerely about his appreciation and I just started thinking "what am I doing - can I really sell this thing?"

What those experiences in August did was help me crystallise priorities. I was still determined to see the process through, but I gained a greater appreciation for just how important it was to find the *right* organisation. I left Vegas feeling like HIBP was much bigger than just me.

Non-binding Bids

Of the 43 companies that received the IM, a subset of those then submitted non-binding bids which essentially means "here's the deal we'd like to do, but there's a heap of due diligence we need to do yet before making a binding commitment". I'm going to be a little vague on that number as I honestly can't remember what I represented to each of these organisations in terms of levels of interest due to the way the bids trickled in. There was a very clear timeline to submit bids given to each potential suitor, but many of them missed it not just by hours or days but in some cases, even weeks. Unsurprisingly, some organisations elected not to submit bids at all and that was really the aim of the IM; to filter out those who were serious with proceeding from those who wouldn't ultimately be suitable. We *had* to ensure the 43 who received the IM

was significantly chopped down.

The non-binding bids were the first time we started to get a true sense of how the various organisations valued the service. It wasn't just the headline value either, it was how much of it was comprised of cash versus equity and over what period of time it would be paid, which brings me to a really key factor in all of this: golden handcuffs. A consistent theme across all the bidding companies was that they wanted me locked in for years and if I changed my mind part way through, I'd pay for it *big time*. I expected that - it wasn't news to me - but I'd be lying if I said it didn't worry me once I started seeing it in writing. If I entered into one of these agreements then, for example, decided I didn't like a strategic change in direction the organisation took and decided to leave, I'd no longer have HIBP, I wouldn't be able to do anything similar for years due to noncompete clauses and I'd be financially penalised *massively*. That weighed more and more heavily on me as things progressed.

The non-binding bids helped us further chop down the list of suitors. There was a massive spread of valuations. Some companies wanted me to perform roles I wasn't comfortable with. Some wanted me to permanently relocate overseas. I hadn't ruled out relocation at the beginning of the process, but there were enough organisations happy for me to be anywhere that it left plenty of options open without giving up my Gold Coast lifestyle (seriously, just look at this place!)

Exclusivity: Then There Was One

Apparently, the way these M&A processes run is that as you really get down to the wire with the final bidders, eventually someone will ask for exclusivity. This grants them a window of time in which they can do extensive due diligence to the exclusion of all other bidders. That might sound a bit selfish on the face of it, but as I'd soon learn this can be a *very* laborious, drawn out and expensive

process. A bidder who believes they're in with a good shot wants to make sure they can make that investment and have a high likelihood of coming out as the victor.

And so in September, we granted exclusivity to a bidder. Now, I'm going to be extra careful here with the words I use because even though there wasn't ultimately a sale, I signed off on all sorts of confidentially terms which prohibit me from sharing anything that might indicate who this bidder was, how much the bid was for or what the terms of the bid were. I hate to be vague (I'm usually super transparent on these things), but I'd also hate to disrespect the privacy of this organisation or land myself in hot water legally. What I will say is that it was a company that met all my criteria both as outlined in the original Project Svalbard post and so far in this one. It was a company I respected and one I had confidence would help me take HIBP in the right direction.

And so began the extensive due diligence. KPMG had warned me about this phase right at the beginning of the process and from memory, the word they used was something akin to "onerous". Let me try and give you a sense of just how true that word was by way of examples and I'm going to pick a handful not just from the company that had exclusivity, but from some of the earlier 43 as well. Among literally thousands of other requests (seriously - the total number was four figures), I was asked for:

- 1. Minutes of all meetings of the board (remember, HIBP is a one-man show)
- 2.Documented processes for when a mobile device reaches end of life (uh... I factory reset it and give it to the kids?)
- 3."Documentation of the Company's technical operations, including but not limited to platform capabilities, database servers, data center operations, network infrastructure, IT policies, SLA's provided to customers, back-up/redundancy plans, and emergency/disaster recovery procedures"

I copied and pasted that last point verbatim - can you imagine how much information needs to go into a response to a question like that?! How HIBP runs across the various Azure services, the Cloudflare dependencies, how I recover if things go wrong and then how that's managed across different autonomous parts of the project such as the HIBP website, the Pwned Passwords service etc etc. This just isn't the sort of stuff you document in a pet project so everything had to be done from scratch.

What I was being asked for during this extensive due diligence phase wasn't coming from the folks I'd initially spoken with in the lead up to their non-binding bid, rather from the leagues of business development and legal folks behind them that needed to get involved in this process. They didn't know who I was, had likely never heard of Have I Been Pwned before this exercise and if I was to take a guess, wouldn't have even known how to pronounce it. But M&A was what they did and they were simply asking all the sorts of questions they would in any other M&A process so I can't begrudge them that. Problem is, it's one thing to get hit with those questions when you're part of a team of people, but it's a whole different thing when you're one bloke on his own.

I don't think I'll ever be able to sufficiently explain all the emotions I felt during this phase of the process. It was an endless series of questions, meetings and if I'm honest, frustration. At one stage, I sat between lawyers arguing backwards and forwards as to whether or not I was a sophisticated investor up to speed with American Securities and Exchange Commission law and if I wasn't, "the deal's off". I got a bill for that argument.

This went backwards and forwards for *months*. Every time we thought the whole thing was done there'd be more questions. More delays. Until it was over.

Then There Was... Zero

The news came very recently. Keeping in mind my previous point regarding

confidentiality and choosing my words carefully, the circumstances that took the bidder out of the running was firstly, entirely unforeseen by the KPMG folks and myself and secondly, in no way related to the HIBP acquisition. It was a change in business model that not only made the deal infeasible from their perspective, but also from mine; some of the most important criteria for the possible suitor were simply no longer there. Collectively, we agreed to put pens down.

After many months of exclusivity with a single organisation and going through crazy amounts of due diligence, the effort involved in scrolling back to the September time frame and starting it all again with another organisation would have been enormous. I also didn't want a situation where I compromised my own principles; the organisation we'd identified as the best possible fit was precisely that - the best possible fit - and all other candidates would mean making concessions I simply couldn't justify. Besides, as this exercise had already demonstrated, there are absolutely no guarantees in this process and going back to square one could very easily result in many more months of effort and no outcome to show for it.

So we wrapped it up, I got the single largest bill I've ever received in my life and then I sat down and started writing this blog post. In doing so, I stopped for the first time since April 2019 and reflected on how much had happened during the process.

A Lot Happens in 11 Months

I onboarded 5 new governments onto HIBP: Austria, Ireland, Norway, Switzerland and Denmark (and a 6th one about to be announced any day now). I loaded 77 new data breaches comprising of 1.7B records into HIBP and signed up almost 400k more individual subscribers to the service. I built and launched the authenticated API and payment process (I really should have done this earlier, I'm so happy with it!)

On a more personal note, I joined the likes of Bruce Schneier, Eugene Kaspersky and Alan Turing (*Alan Turing*!!) in the Infosecurity Hall of Fame. I spoke at CERN. I visited 2 new countries for the first time (Israel and Hungary) and keynoted events there, plus a heap of talks in more familiar places, a bunch of workshops, I still wrote blog posts and somehow - miraculously - never missed a weekly video.

On the M&A front, I had to learn about normalised EBITDA, revenue multiples and ARR. I met literally *hundreds* of people in person regarding Project Svalbard during both the San Francisco meetings and travel to other parts of the US and the world.

During all of this, I still had to run HIBP in a "business as usual" fashion. I still manually verified every breach, hand edited every logo of a pwned company, issued (and chased) every invoice, did the tax returns and prepared the <u>business</u> activity statements. In other words, all the stuff I'd always done for years still had to be done regardless of how menial it was, none of that went away. I'm detailing all of this here to help explain what I need to do next...

So. What's Next for HIBP and for Me?

To be honest, I need some time to recover. What I've explained in this post will never adequately illustrate just how stressful this process was. I need some time where I'm not waking up dreading how much work will have landed in my inbox overnight. I need some time to write more code and more blog posts, two things that remain my passion but had to take a back seat during this process. I'll still keep running HIBP as I always have, but I need the head-space to get my energy levels back up and plan the next phase. I've (almost entirely) cleared my calendar for the next few months to give me that much-needed time out and with coronavirus causing a heap of conferences to be cancelled and travel plans to be disrupted, it's probably not a bad time to stay home anyway.

Having said that, there are things that have become abundantly clear during the M&A process that I'm confident will feature in that next phase. I need more support, for one. I can't be the single person responsible for everything so I'll be considering the best way to start delegating workload. That'll not only help me run the service as it stands today, but it'll help me expand it to do so many of the things I'd wanted to in a post-acquisition world. It'll also allow me to work towards no longer being the single point of failure; there has to be a contingency plan for if I get taken down in a freak drop bear accident.

One of the things I'm *really* excited about is a concept I've had bubbling away in the back of my mind for a couple of years now about how the industry as a whole can better tackle the flood of data breaches we're seeing. I floated this idea past each of the companies I met with during Project Svalbard and the support for it was overwhelming, even from those organisations that knew very early on they wouldn't be bidding. For those reading this that were part of those discussions, I'm determined to make it happen this year and I *will* be in touch!

Another area I expect to focus on a lot more is to leverage the more formal relationships I established during the process with governments, regulators and law enforcement. It's an interesting time right now where there's clearly a lot of support for HIBP and the way it operates, but also a lot of focus on privacy and people having control of their own data which poses some interesting challenges. Here's a simple example of the paradox I want to tackle with these groups: we all want privacy but we also all want to know where and how our data has been exposed and *what* the data is, how do we achieve both objectives? It's non-trivial for many, many reasons, but it's also important and HIBP has a role to play in the solution.

The list of all the things I want to tackle in the post-Svalbard era is lengthy and that's also why I need the downtime: to be able to focus, prioritise and take HIBP forward with more enthusiasm and energy than ever.

Summary

I saw <u>a comment only last week on my traffic spike blog post</u> that really brought home how much I love running this project:

You're living the dream and you make it look good Troy.

I love what I do. This is a fascinating industry that continues to challenge me in all sorts of ways I never expected and there's not been a moment where I've felt bored or uninspired by it. To be able to continue running HIBP and shepherding it forward remains the dream, regardless of who owns it. So, I'll finish this blog post on the same note I finished the last Project Svalbard one:

I've made this decision at a time where I have complete control of the process.

And so it remains today and for the foreseeable future, with HIBP as an independently operating service designed to do good after bad things happen. Thank you for reading this far, thank you for supporting both HIBP and myself, I'm off to have that board meeting 5#





There's no place like home 💚 🔠 🐨 🐛 📥 💥 💫















12:02 PM · Feb 27, 2020 from Narrowneck Beach

Comments

Firstly, let me say how much we appreciate your hard work! The company I work for recently went through the same Due Diligence process and it took a team of people many, many hours to produce the same kind of documentation you speak of, so for you to do that as a one man band, I shudder to imagine the stress of it all. Bravo.

On the plus side, you now have the basis for making HaveIBeenPwned an actual company as you have your processes written down and documented in a fashion that Lawyers (as well as any new staff) can easily understand.

In your post you speak about working with the IT security/wider internet industry on improving the trust relationship between users/companies and I wholeheartedly agree that this can only be a good thing.

Are you able to accept donations? (I know you had a "buy me a beer" link ages ago, is that still a thing?)

Troy: Beers are still very much welcomed <a>https://haveibeenpwned.com/...

Knowing what a pain selling a business is like (my wife recently sold one of hers), I can commiserate.

Have you considered starting a company/foundation to continue to run HIBP as you envisage, and employing someone to do the donkey work, while you focus on the things you enjoy?

Troy: HIBP already runs under a company and yes, employing people is certainly under consideration, I just need some time to work out the best possible path forward.

Considering a lot of governments (and more to come?) uses Have I Been Pwned, as well as organizations. Would it make sense to go in the direction of efforts like Mozilla? Receive funding/donations and uses that to run the project(s)?

If a thousand companies and organizations and governments just donated/sponsored a few bucks each (per month) that would go some way towards hiring some part time help maybe? This would also ensure the whole thing remains independent and unbiased insofar that any economic interests is the self sustainability of the project(s).

For example, Google is sponsoring Mozilla, which I find very noble.

Another interesting example is VirusTotal https://www.virustotal.com/

Check out their "How It Works" in the bottom left of the footer. The project was started by Google if I recall correctly.

They have free virus scanning against all major (cranked to max paranoia) virus scanners. The more advanced features (which has been added over time) is non-free.

Another example would be <u>Gitlab.com</u> <u>https://about.gitlab.com/co...</u>

These are partly altruistic/partly capitalistic in the way they do things (and it works). No features are removed (sometimes new features are added or unlocked for free use too).

It should be pretty obvious that what you (Troy Hunt) have done for Project Svalbard up to now has worked and you may be one of very few that know how to keep the ship on course.

Regardless of what you end up doing going forward, I'll be cheering for you!

Troy: There's a bunch of different ways of tackling this (and they're just the ones I've thought of already!) and yes, I do see funding and / or donations possibly featuring in there in the future. All good suggestions you've made there Roger, thank you for taking the time to write it \bigcirc

Epilogue

Conservatively, Project Svalbard cost me about A\$400k. That's just in bills I paid primarily to KPMG, add on travel, accommodation and if you really want to get a true sense of the cost, opportunity loss while I gave this process my all, and the number is a hell of a lot higher again. I don't regret it though, rather I look on it as one of life's lessons and as I've said many times before, I'm happy with where I am in life today and I'm only here because I've been defined by my experiences, this one included.

The deal falling over gave me a chance to stabilise my life, especially as it related to my new relationship and the time I spent with my family. COVID-19, if I'm purely selfish, was a godsend; no more travel, no more pressure to be somewhere else and no questions raised about my ability to be around for the kids. It got us into a routine the likes I'd never had before with regular exercise, time with my parents and building out some lovely relationships with neighbours. I closed my exercise rings every single day, often walking to the beach or bike riding, hitting the wakeboard park and going for bushwalks.

What would life had looked like had the HIBP sale gone through and there was no COVID-19? By comparison - terrible. I can't think of a more appropriate word. The deal died of natural causes, and I couldn't have been happier.

SUSTAINING PERFORMANCE UNDER EXTREME STRESS

I've always been resilient. If I'm stuck on a problem then I just work away at it until I solve it, even if it takes way too long. If I'm exercising then I keep going until I reach my goal, regardless of the pain. And now I found myself dealing with something on an emotional level I'd never experienced before, but with no option other than to push through it.

As I say at the very start of this post, this was a piece I started writing at the absolute zenith of stressful times. It wasn't just the breakdown of a 20-year relationship nor was it the sale process I was going through with HIBP, a lot of the stress was dealing with the constant oscillation of emotions. On the one hand, I was mourning the loss of the idea of marriage (for a non-traditional guy, I struggled a lot more with the stigma of divorce than what I probably should have), yet on the other hand I was at the very beginning of a new relationship which was an exciting time. I was tired and run down and missed my kids (I've already mentioned the 11 week trip), yet I was flying in style and wining and dining. The lows were the lowest and the highs were the highest and that's what really sticks in my mind about that period: oscillating emotions.

I published this post once I was finally ready to talk about my divorce. Some people had worked it out already (I remember someone noticing I no longer wore my wedding ring in my weekly videos well before my ex-wife even noticed), but I was conscious I lived a pretty public life and not talking about this just felt like a glaring omission in my usual transparency. On reflection, I think what it was that made me finally ready to publish this post was having turned the corner of getting over the relationship breakdown. For all sorts of reasons, I was well and truly done worrying about it and just wanted to put the whole thing behind me.

02 JULY 2020

September. It was at the absolute zenith of stress; a time when I had never been under as much pressure as I was right at that moment. Project Svalbard (the sale of HIBP which ultimately turned out to be a no-sale) was a huge part of that and it was all happening whilst still being solely responsible for running the project. That much was very broadly known publicly, but what I haven't spoken about until now is that earlier that year, my wife and I had decided to separate and later divorce. As part of attempting to rebuild my life, I was also in the midst of buying another house, a stressful process at the best of time let alone under these circumstances whilst on the other side of the world. It was extreme stress the likes I'd never dealt with before at a time when the demands on me were at an all-time high, so I started writing this blog post, adding to it at the worst of times. Here's how I sustained my performance whilst under extreme stress:

I Leaned on Friends More Than I Ever Had Before

I realised something very profound last year; I've very rarely discussed my emotional state with friends. Maybe that's "just what blokes do" (or don't do), but it certainly wasn't a conscious decision on my behalf. It wasn't until the stress really started mounting early last year that I actually made a conscious effort to do this. Putting it in words now seems almost stupidly obvious, but there's a lot of evidence around the benefits of friendship on mental health:

It can be hard to talk to family members about mental health. That's why it's important to have healthy friendships to turn to in times of need. Our

friends can be that ear to talk to, shoulder to lean on and nonjudgmental perspective that we need. They can also help increase our sense of belonging, improve our self-confidence and help reduce stress and anxiety.

Last year and early this year, it meant spending a bunch of time with friends in person during my travels. Since Feb this year as travel has become a thing of the past, it's meant talking to friends in different parts of the world every couple of days. Often those discussions have directly focused on the stresses in life but equally often, they've been an opportunity to bond around less contentious common interests; cars, tech, family. The quote above about helping to increase a sense of belonging really nails it.

The thing that perhaps surprised me most about those discussions with friends was how much their own stories resonated with mine. I mean that across all the fronts I was feeling the stress on too; whilst in San Francisco in particular, I spent a bunch of time with people I knew well who'd been through similar business processes and as for the things stressing me in my personal life, it felt like every second person I confided in had a similar story. Finding common ground with friends was always a huge relief; I wasn't alone in what I was going through.

I Did My Utmost to Not Make Decisions Based on Emotions

Emotions have been high during this period, both professionally and personally. More than anything, it was the *unpredictability* of emotions that got me; I could be cruising along thinking everything was on track then wammo! An email, a text message or a phone call would suddenly throw everything back into turmoil. I'd be upset. Angry. Vengeful. But none of these feelings would help me make rational decisions.

Frequently, I'd simply sit on an email for a day. I'd sit on *my own* emails for a day, granting time to reflect on whether my words represented the best path forward or merely reflected my emotional state at the time. A perfect example is that the house purchase fell through due to the vendor not being agreeable to the terms I set forth. I received their reply and was initially upset. I sat on the email, went and did a conference talk, drank some beer, had a sleep and responded the next day, cancelling the deal. It hurt to do that because I *really* wanted the house, but I also knew that "want" wasn't enough, it had to actually make sense and without agreeing to my terms, it simply didn't.

I haven't always gotten this right and there hasn't always been the luxury of time between emotion and response, but as a strategy to keep peace and maintain sanity, it's proven invaluable time and time again. I can't think of a time where I slept on a response and didn't tone it down a bit.

I Stayed Focused on The Bigger Picture (and the Small Steps That Would Take Me Towards it)

I was always looking a year or more ahead and I had a very clear picture in my mind of how I wanted my life to look like in the future. Stress has a way of clouding judgement and causing you to make irrational decisions, many of which might feel right at the time, but don't ultimately further your life goals. I had a vision of what my future would look like (and obviously given the HIBP no-sale, reality hasn't always aligned with the vision), and everything that was happening as I wrote this blog post had to support that objective. But there were also *massive* changes in my life that had to be dealt with here and now, and there was only one way to do it:

I like <u>the way Psychology Today explains this adage</u>, by breaking those steps down into goals that must be:

- 1.**S**pecific
- 2.Measurable
- 3. Attainable
- 4.Relevant
- 5.Time-bound

Consider what was required to achieve the big picture goals I had; everything from literally hundreds of meetings, thousands of emails, endless proposals, terms sheets, negotiations - and that was just on the HIBP front. Throw in the stress, emotion and frankly, some pretty dark moments on the relationship side of things and consider how totally overwhelming it can all feel.

I tackled it by focusing on the very next thing I needed to do to; the single, *attainable* thing I could do to move me towards a goal. Complete some financial documents. Schedule a meeting. Agree on some key deliverables. So long as the activity was an enabler of that big picture it didn't matter that it was a little thing, it was *progress*.

I Tried Not to Sweat the Small Stuff

It's *so* easy to get bogged down in detail and derailed from focusing on what's actually important, that <u>there's literally a book on it</u>:

#1NEW YORK TIMES BESTSELLER

DON'T SWEAT THE SMALL STUFF...

and it's all small stuff



SIMPLE WAYS TO KEEP THE LITTLE THINGS FROM TAKING

OVER YOUR LIFE

RICHARD CARLSON, PH.D.

COAUTHOR OF HANDBOOK FOR THE SOUL

I can think of many occasions across all the various things that put me under stress this last year and a half where I literally concluded "fuck it - it just doesn't matter enough". They were things that by any reasonable measure I had every right to be upset about, but equally they were things that *had* I gotten upset about them, they'd derail me from focusing on that bigger picture.

Legal jargon in contracts is a prime example. I recall one occasion where lawyers on my side of the HIBP deal were arguing with lawyers on the other side about whether or not I was a "sophisticated investor". I needed to be in order to receive the proposed equity component and unless we agreed that I was, the exact words I heard were "the deal's off". It was an obnoxious comment about a ridiculous premise, but ultimately, we concluded that the real world impact of the clause was likely negligible and further arguing about it really didn't serve my own purposes.

I Moved on Quickly from Setbacks

There were so many outcomes along the way that frankly, felt devastating. Incidents and events that left me fuming, emotional and sometimes, pretty inconsolable. It was so easy for these things to eat me up and consume me, taking my focus away from that big picture and keeping me from moving forward towards that bigger goal.

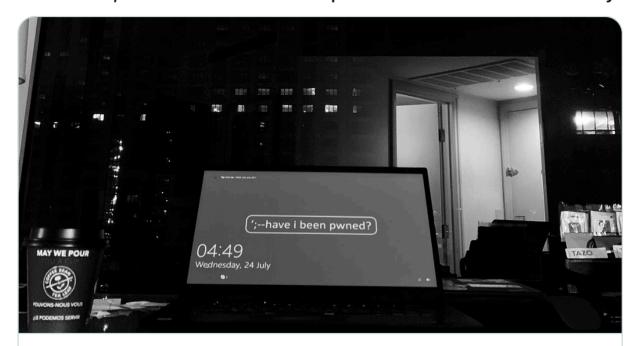
I found I kept going through the same cycle after a setback and it tracked pretty closely to the whole <u>Kübler-Ross model of 5 stages of grief</u>. I'd very quickly move through denial and anger, blast through bargaining and depression and get to acceptance. I tried hard to bring myself to that last stage and I remember thinking so many times on the way there "this feels much worse now than it will tomorrow or the next day".

In thinking of an example to illustrate this, the following tweet and excerpt from the "no sale" blog post came immediately to mind:





Project Svalbard was the initiative to find a new home for @haveibeenpwned. After 11 months, the project has now run its course; HIBP will remain independent. Here's the full story:



Project Svalbard, Have I Been Pwned and its Ongoing Independence This is going to be a lengthy blog post so let me use this opening paragraph as a summary of where Project Svalbard is at: Have I Been ... \mathscr{S} troyhunt.com

10:07 AM · Mar 2, 2020



According to the lock screen, I took the photo below at 04:49 on the 24th of July last year. I was in yet another bland, nondescript hotel room, drinking bad coffee in an attempt to stave off the jet lag. I'd arrived in San Francisco a few days earlier after barely making my connection in Helsinki, literally running through the airport. My bag hadn't made it. I was tired, alone, emotional and if I'm honest, at an all-time low.

I felt like shit at that moment, but it was temporary and I had *just* enough sanity left to know that the feeling would pass. Just. But it always *did* pass and there'd be something else of a much more positive nature happen the very next day.

I Always Thought 3 Steps Ahead

Let me begin by saying this: I didn't always get this right (far from it) and on multiple occasions I got blindsided by things I never saw coming (the circumstances under which HIBP ultimately didn't sell is a perfect example). But the basic premise is that before expressing my position on something, I'd consider the range of possible responses I'd receive. Let's say there were 3 of them; for each of those 3 possible responses I'd not only consider how I'd respond to each, but how each of my responses would then be received. Same again for how I'd respond to each of those and in my mind, I was drawing out a mental image of 3 ^ 3 different possible outcomes - which one did I want? It was an exercise that enabled me to look much further down the road and consider whether it aligned to an earlier point in this blog post - my big picture.

This requires time, practice and patience and as I said in the opening, I didn't always get this right. You can't always be aware of all the factors influencing third parties nor can you be aware of all the cards they hold, but without doubt, this way of approaching any negotiation is enormously valuable. It also forced me to empathise; how will other parties feel? What's the most natural reaction they'll then have?

In my mind, this is akin to a "choose your own adventure" book; at each crossroad there are different ways you can go. Each of those then has their own crossroad as do those ones too. Before making a decision at that first intersection, I want to know what the next 3 will look like.

I Drank Beer

Treat this less as a suggestion to consume alcohol and more as a representation of taking time out for yourself. For me, having a beer is something I associate with switching off from the everyday stresses. I very rarely drink alcohol when working (now coffee, that's another story!) and treat beer as an opportunity to "down tools" and relax.

I drank beer on my own in a pub:





Replying to @gheja_

Cheers!



5:27 AM · Sep 24, 2019

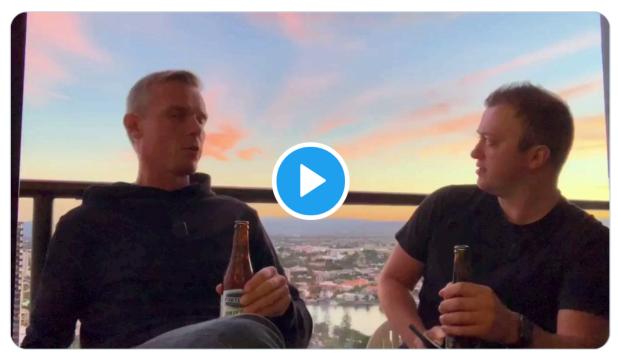
(j)

I drank beer with friends:





Weekly video with @Scott_Helme done! I'll edit and publish tomorrow, just a quick teaser for now - check that Gold Coast sunset



11:23 PM · May 30, 2019 from Surfers Paradise Marriott Resort & Spa

(i)

I found new ways to request beer:





Merry X'mas! 🎄



6:55 AM · Dec 24, 2019

(i)

The point is that I made a conscious effort most days to tune out and give my brain a rest. A good mate of mine is convinced meditation is an equal of beer in terms of helping him disengage from daily life and maybe he's right, I just don't have the patience for it (yet). Find your beer, whether it be actual beer or an activity which allows you to do what the process of going and having a cold one does for me.

I Threw Myself into Exercise and Health

From beer to physical wellbeing: I was trying to find a tweet to illustrate the point, and this one nails it:





Still doing the #AusCERT2019 Q&A while eating poke before hitting the tennis court. 15 mins to go!



4:47 PM · May 20, 2019



At this time, I was now well into Project Svalbard, I'd separated from my wife and per the caption, I was preparing to deliver a keynote at Australia's premier security conference. When I first started to really feel the stress, I absolutely

threw myself into exercise:



Gav is both my son Ari's and my own tennis coach. I literally said "Gav, book me in every day at the hottest possible time" and when the weekend came, I'd play with Ari as well. The standing commitment each day forced me to get out on the court and focus on something other than life's stresses.

Per the earlier image, I was also getting right into <u>Poké Bowls</u> which meant a lot of raw fish, brown rice and greens like edamame and seaweed. I'd order it on Uber Eats, it'd arrive at my door and IMHO, it's genuinely delicious. Physical health has a profound effect on your ability to perform mentally, particularly when you're under extreme stress. Exercise in particular has <u>very well-documented benefits when it comes to depression, anxiety and stress</u>.

Despite the emotional turmoil of recent times, I'm in great shape physically with a typical week including running, bike riding, tennis and wake boarding. I'm about to pass 3 months of closing all rings on the Apple watch every single day (amazing how much not travelling helps you do that!) and I can really see those benefits showing in the kids too when they share the activities with me.

I Established Stable Routines

In a tumultuous period like this, it's easy for routine to go out the window. Most people have some form of routine which establishes consistency in their life, for example going to work each day. A regular social commitment. A Sunday roast dinner. I spent 243 days travelling last year so consistency was near non-

existent.

A saving grace has been <u>my weekly update videos</u>. Every single week, without fail, I've done the video. Sometimes they've been at the worst of times, needing to record and put my face in front of the world after feeling emotional / jetlagged / broken (and a big shout out to those who commented to that effect!) But what those videos did was give me a small sliver of consistent predictability in life. During each week I'd take notes on content, pull myself together then sit down and record.

Same again for my blogging and drafting this one in particular was a big part of that. For the last 11 years, I've written about most of the things in my life that have been important. Writing transparently about what's going on in my life has become a part of my routine and indeed, a part of my identity. It feel "off brand", for want of a better term, when I don't.

As things have stabilised this year, I've been able to broaden those routines with regular tennis, time with my family and simply walking down to the beach most mornings:





A picture alone can't show just how epicly beautiful it is here today 🔆 🥗 🥣



11:59 AM · Jun 21, 2020 from Main Beach, Gold Coast



I snapped that pic last week after watching a humpback whale and her calf cruising by, probably just 50m offshore. It was a moment of reflection following a period of great turmoil; it's been both the highest of highs and lowest of lows. But now, being at home and finally having stability it's crystal clear: this is a routine that's going to stick around for the long term.

Closing

This was a heartfelt blog post about some momentous events in my life. By all means, please comment, share your experiences and ask questions but avoid topics related to my relationship. As much as I'm open about the emotions I went through and how I dealt with them, details of a personal nature are something that will remain that way. Thank you.

Comments

Thanks for sharing. My own personal moment of highest stress came in 2017 after my wife left overnight, the weekend before I was on a fortnight of early cover at the office, and forced me to move across the country (one carload at a time), whilst I figured out childcare with my family, and had to deal with rebuilding my life whilst commuting, moving out of my sister's spare bedroom, sorting out a school for my son, moving all of my stuff one car load at a time, a tax return, Christmas (!), and closing on a house I was contractually obliged to buy but no longer wanted.

At some point you just have to accept that you're spinning so many plates, you're gonna drop a few of them. I forgot to go to a Christening!

Eventually I got organized, and found some CBT-based therapies that helped me organise my brain and deal with everything. I put down the Machine Learning course I was trying to do in my spare time; very interesting but just not important at the time.

I started dating again. I found someone I've known for 17 years. I proposed. We're getting married in October.

We moved into the damn house 🥩

I'm pleased to hear you seem to be finding your way through a very tough time, and I'm not surprised it took you nearly a year to write this.

I keep my 2018 diary as a reminder of what a crazy time it was, and also to tell myself that I am so much more capable than I sometimes believe.

Thanks Troy for sharing. Didn't realise you were going through a divorce amongst all the other stuff. While it's still raw now, it will get better over time. The saying goes - as one door closes, another door opens. I literally wouldn't of been born if my dad's first marriage didn't end and he never met my mother and re-married. This fact helps put a lot of my own personal or professional setbacks in perspective.

Good on ya for posting this.

Some of my experiences / notes to my younger self and perhaps some words of encouragement in a difficult time.

Adversity builds character. You've shown character in the way you dealt with what must have been a very difficult time.

In difficult times we often want answers or confirmation that the next steps we're taking are the right ones. Of course there's usually no answers in times like these.

One of the benefits of growing older is that you can draw on past experiences, knowing that you've dealt with difficult situations before and came out on top. Drawing on that knowledge has often pushed me through difficult times.

Simply getting up each day and putting one foot in front of the next, keeping busy and continuing with routines and exercise is remarkably effective. You've articulated it pretty well: "Routines bring consistency and predictability". In difficult times it's also an area of your life you are in control of when everything else is seemingly out of control.

It was helpful to read how you dealt with this. I distinctly remember dealing with a crossroad in my life, thinking where to from here? How do I deal with this and move forward? How do other people deal with it?

A word of encouragement. In times of difficulty it's often difficult to see the way out. The

most difficult periods in my life have always opened new doors and looking back I never look back with regret or sadness on difficult periods. And that's how the balance of life works.

Just keep walking mate, she'll be right.

Troy: Like the bit about difficult periods opening new doors, this whole situation combined with COVID as well has opened a bunch. I spend a lot more time at the beach. I talk to neighbours I never would have otherwise. I got a new bike and ride a lot. There are certainly upsides if you choose to find them.

Epilogue

I wasn't as candid in the blog post as what I feel I can be here, well over a year later; there were times where I was an absolute emotional wreck. One of the lowest points I can remember in the entire experience happened the night before I snapped the photo of my laptop at 04:49 embedded in this post. I relayed that story I'd previously shared in the Svalbard post about having arrived in San Francisco jet lagged, alone, emotional and with no bags, but I didn't explain what "emotional" looked like. I'd called my parents the night before and just been completely inconsolable, unable to speak through the tears. I couldn't get a word out and there was nothing they could say to help the situation. It must have been terrible for them being on the other side of the world and hearing their son in that state, unable to do anything about it.

I picked myself up the next day and went and bought a change of clothes. My bag arrived, as did one of the KPMG guys and shortly thereafter, some cold beers and good-hearted chats. Days later, Charlotte arrived too and whilst none of the stresses that had upset me so much earlier on had gone, it felt like the support team had arrived and I could get back to focusing on the job I was there to do. Usually when something is stressing me, I reach a pretty early realisation that this is something that will pass and whilst this episode

demonstrated precisely that, it really knocked me about emotionally.

What I didn't expect from this post was the extent of the outpouring of support that followed it. There were some lovely comments on the piece itself, but the personal messages went way beyond that. I received a heap of emails and DMs from people both sharing their own stories and supporting me as I told mine. It made a genuine difference; I mentioned "stigma" in the intro, and it was only once I understood how many other people were going through much of the same stuff I was did I really come to terms with my own situation.

I realised after posting this piece that it was actually helping other people deal with their own stigmas, especially around mental health and especially as it relates to men being willing to talk about it. I didn't give it much thought at the time as I never considered what I was going through to be a mental health issue, but it evidently touched a nerve with many others who in all likelihood, were dealing with much more serious issues than I was. With this in mind, I started drafting another post. This time it's specifically about dealing with divorce and like this one here, it's a work in progress I keep adding to. One day, when I finally feel the time is right, I'll publish it.

HACKING GRINDR ACCOUNTS WITH COPY AND PASTE

I think all of us have a little bit of the hacker ethos inside of us. There's something that makes us want to understand how things work, then work out how to break them. I reckon we all started out that way as kids; curious, inquisitive little beings that would just as soon pull a toy apart as play with it in the way the maker intended. Some people grow out of that, others, well, not so much.

Despite how long I've been writing code for and later, doing the infosec thing, I just never grow tired of making software do stuff it was never designed to. Bugs like the one in Grindr are especially interesting because firstly, this one was just ridiculously simple and secondly, because it's Grindr. I'd heard the name before and I knew it was a gay dating service, but I knew little beyond that. I knew it would be a big story because of the nature of the service and that alone really piqued my interest, plus I had an idea of how I could make the story extra entertaining...

03 OCTOBER 2020

Sexuality, relationships and online dating are all rather personal things. They're aspects of our lives that many people choose to keep private or at the very least, share only with people of our choosing. Grindr is "The World's Largest Social Networking App for Gay, Bi, Trans, and Queer People" which for many people, makes it particularly sensitive. It's sensitive not just because by using the site it implies one's sexual orientation, but because of the sometimes severe ramifications of fitting within Grindr's target demographic.

For example, in 2014 Egypt's police were found to be using Grindr to "trap gay people" which was particularly concerning in a country not exactly up to speed with LGBT equality. Another demonstration of how valuable Grindr data is came last year when the US gov deemed that Chinese ownership of the service constituted a national security risk. In short, Grindr data is very personal and inevitably, very sensitive for multiple reasons.

Earlier this week I received a Twitter DM from security researcher <u>Wassime</u> BOUIMADAGHENE:

I contact you because i reported a serious security issue to one of the biggest dating applications for gays (Grindr) but the vendor keep ignoring me !

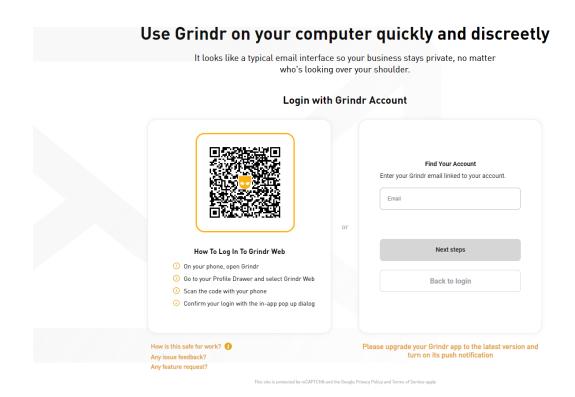
I sent them all the technical details but no way. The vulnerability allow an attacker to hijack any account.

He wanted help in disclosing what he believed was a serious security vulnerability and clearly, he was hitting a brick wall. I asked for technical detail so I could validated the authenticity of his claim and the info duly arrived. On a surface of it, things looked bad: complete account takeover with a very trivial attack. But I wanted to verify the attack and do so without violating anyone's privacy so I asked Scott Helme for support:

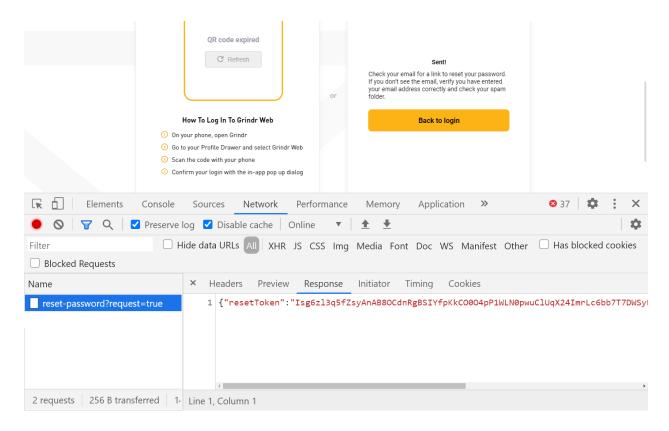


Scott's dealt with plenty of security issues like this in the past, plus he helped me out with the Nissan Leaf disclosure a few years ago too and was happy to help. All I needed was for Scott to create an account and let me know the email address he used which in this case, was test@scotthelme.co.uk.

The account takeover all began with the Grindr password reset page:



I entered Scott's address, solved a Captcha and then received the following response:



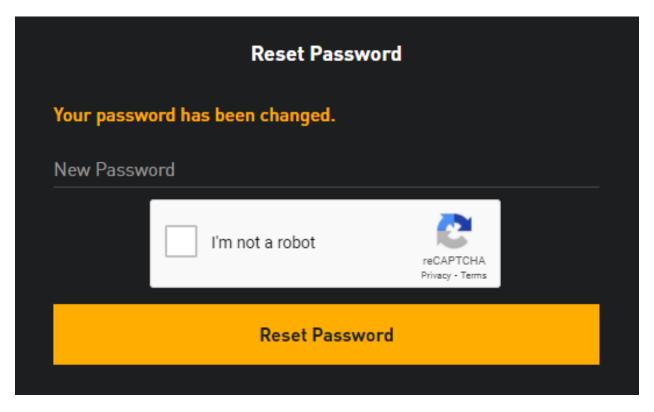
I've popped open the dev tools because the reset token in the response is key. In fact, it's the key and I copied it onto the clipboard before pasting it into the following URL:

 $\label{local-count} https://neo-account.grindr.com/v3/user/password/reset? resetToken=Isg6zl3q5fZsyAnAB80CdnRgBSIYfpKkC0004pP1WLN0pwuClUqX24ImrLc6bb7T7DWSyFMG5lREHQmS4CsFR5uh8GEYQxF6Z6V5hsi3vSTuilXzgKRRItwdDIjmSWdq&email=test@scotthelme.co.uk$

You'll see both the token and Scott's email address in that URL. It's easy for anyone to establish this pattern by creating their own Grindr account then performing a password reset and looking at the contents of the email they receive. When loading that URL, I was prompted to set a new password and pass the Captcha:

Reset Password		
test@scotth	elme.co.uk	
New Passw	ord	
	I'm not a robot	reCAPTCHA Privacy - Terms
	Reset Passwor	-d

And that's it - the password was changed:

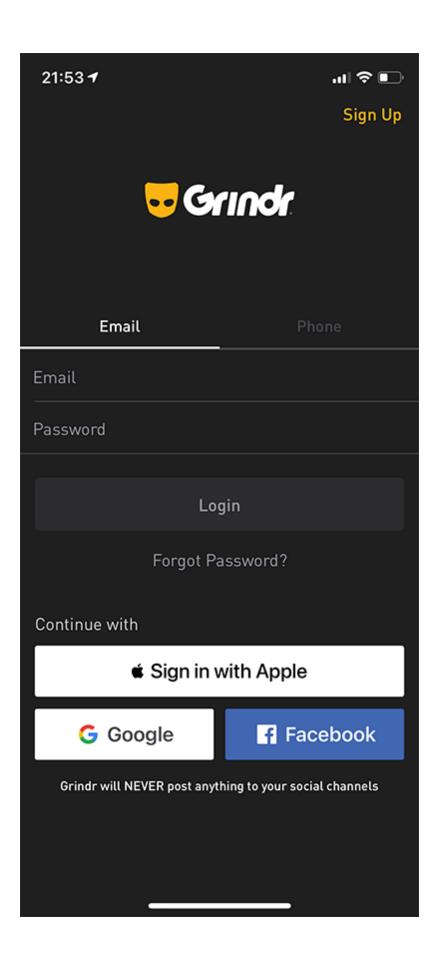


So I logged in to the account but was immediately presented with the following

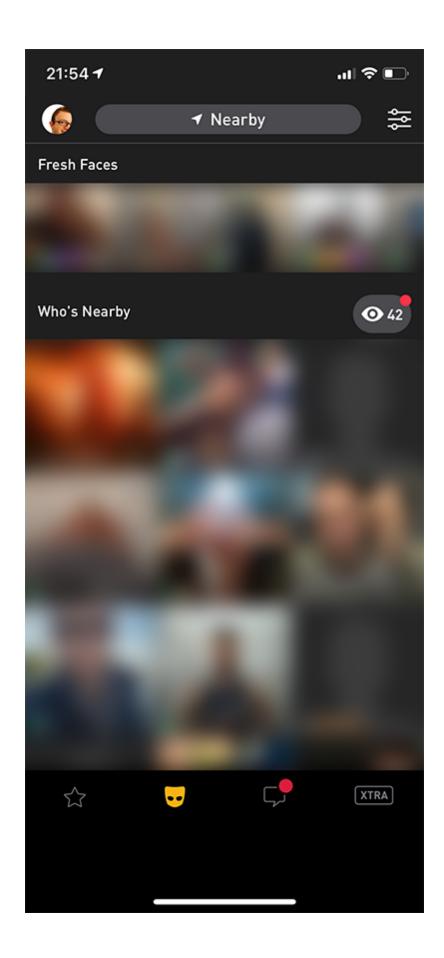
screen:

It looks like a typical email interface so your business stays private, no matter who's looking over your shoulder. Sync Chat? Your browser is attempting to connect to your device. OK For better security, we need you to confirm login in Grindr App. On your phone. Go to your Proff Scan the code v Confirm your to I don't receive notification OK Chass upported wour Grindr and to the latest warsloaded.

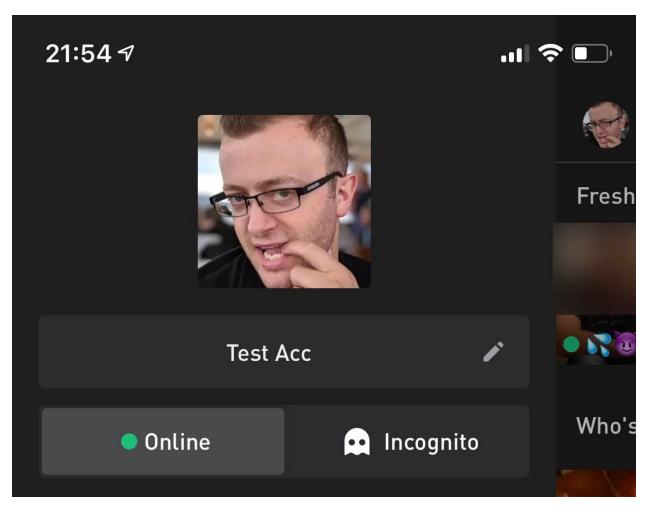
Huh, so you need the app? Alrighty then, let's just log in via the app:



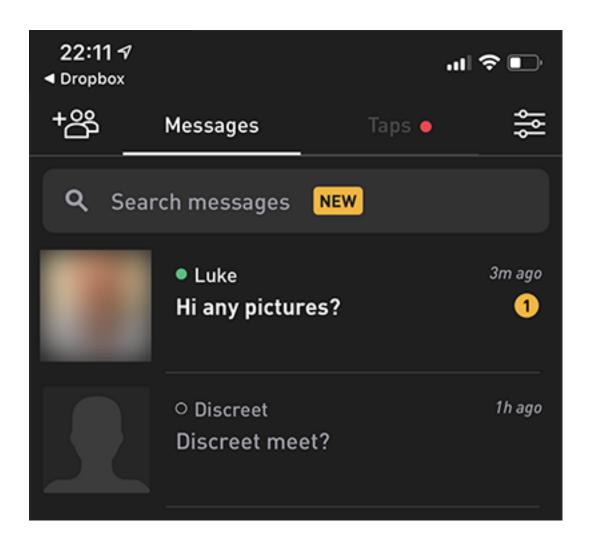
And... I'm in!



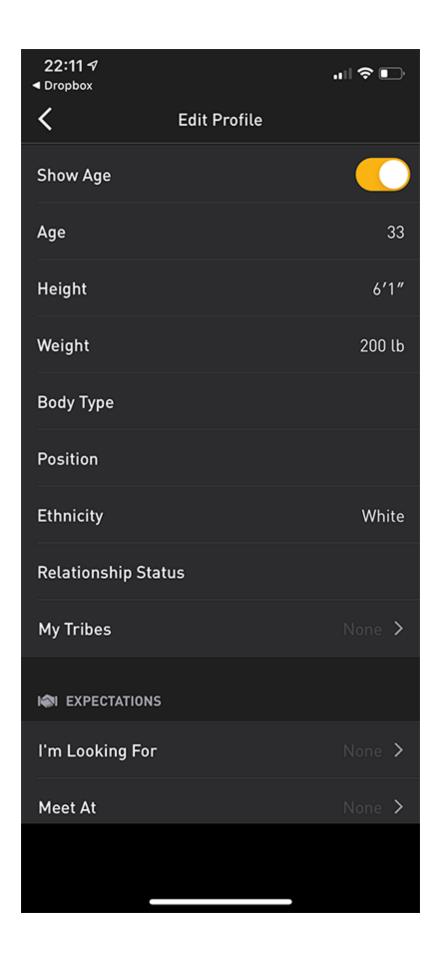
Full account takeover. What that means is access to everything the original Grindr account holder had access to, for example, their profile pic (which I immediately changed to a more appropriate one):

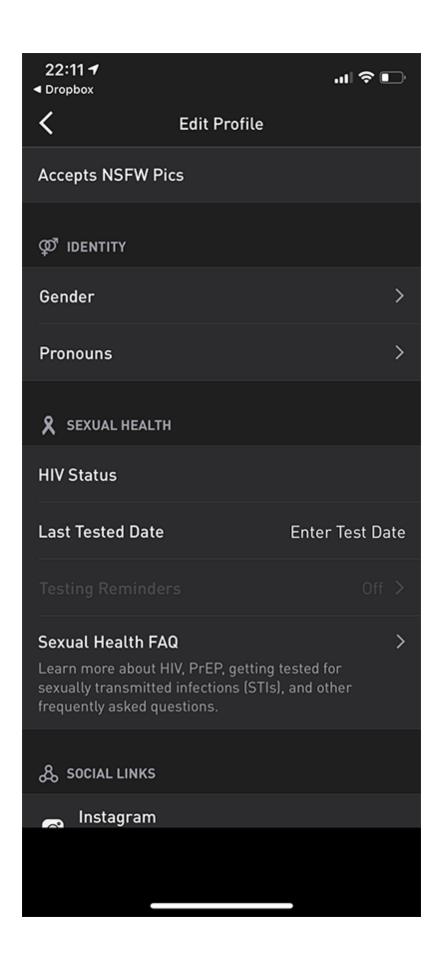


Around this time, Scott started receiving private messages, both a request to meet personally and a request for pics:



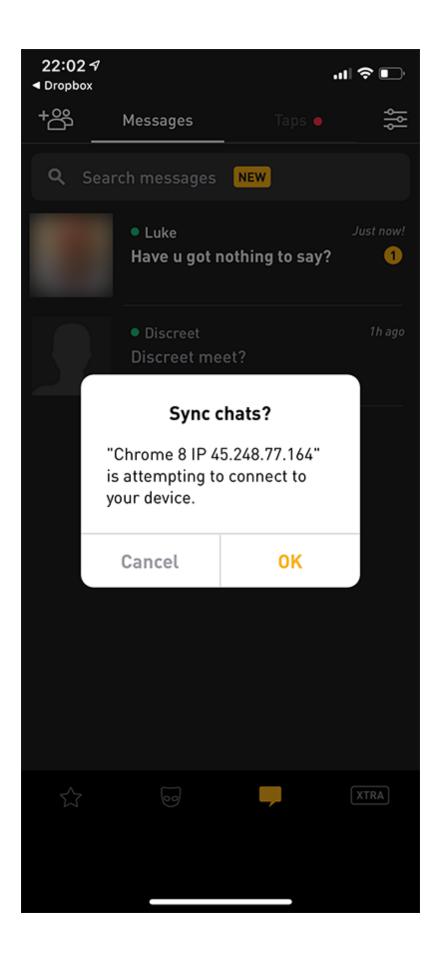
The conversation with Luke went downhill pretty quickly and I can't reproduce it here, but the thought of that dialogue (and if he'd sent them, his pics) being accessed by unknown third parties is extremely concerning. Consider also the extent of personal information Grindr collects and as with Scott's messages, any completed fields here would immediately be on display to anyone who accessed his account simply by knowing his email address:



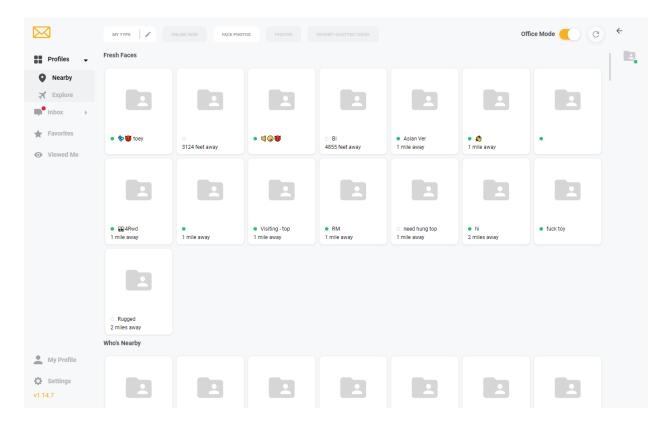


A couple of years ago it made headlines when <u>Grindr was found to be sending HIV status off to third parties</u> and given the sensitivity of this data, rightly so. This, along with many of the other fields above, is what makes it so sensational that the data was so trivially accessible by anyone who could exploit this simple flaw.

And as for the website I couldn't log into without being deferred back to the mobile app? Now that I'd logged into the app with Scott's new password, subsequent attempts simply allowed me to authorise the login request myself:



And that's it - I'm in on the website too:



This is one of the most basic account takeover techniques I've seen. I cannot fathom why the reset token - which should be a secret key - is returned in the response body of an anonymously issued request. The ease of exploit is unbelievably low and the impact is obviously significant, so clearly this is something to be taken seriously...

Except it wasn't. The person who forwarded this vulnerability also shared their chat history with Grindr support. After some to-and-fro, he provided full details sufficient to easily verify the account takeover approach on September 24. The Grindr support rep stated that he had "escalated it to our developers" and immediately flagged the ticket as "resolved". My contact followed up the next day and asked for a status update and got... crickets. The following day, he attempted to contact the help / support email addresses as well and after 5 days of waiting and not receiving a response, contacted me. He also shared a screenshot of his attempt to reach Grindr via Twitter DM which, like the other attempts to report the vulnerability, fell on deaf ears.

So I tried to find a security contact at Grindr myself:



I'm conscious that sending a tweet like that elicits all the sorts of responses that inevitably followed it and implies that something cyber is amiss with Grindr. I only tweet publicly once reasonable attempts to make contact privately fail and based on the previous paragraph, those attempts were more than reasonable. A friend actually DM'd me on Twitter and suggested the following:

Not sure that Grindr tweet was necessary, given their DMs are open and they reached out to you fairly soon after

This is why I didn't DM them:





Hello

i would like to kow if you have a reponsible disclosure policy

i wan to report a high critical security issue, it can be used to takeover the accounts of grindr users

23 sept. 2020 à 10:15 PM ✓

please contact me for more details:

23 sept. 2020 à 10:20 PM ✓

i created a ticket for this security issue

[Grindr Support] Re: #4

23 sept. 2020 à 11:48 PM ✓

please bring it to the attention of the dev team / security team, since it's a serious issue.

23 sept. 2020 à 11:50 PM ✓

That route was tried and failed and I suggest the only reason their Twitter account publicly replied to me was because my tweet garnered a lot of interest.

After my tweet went out. I had multiple people immediately reach out and provide me with contact info for their security team. I forwarded on the original report and within about an hour and a half of the tweet, the vulnerable resource was offline. Shortly after, it came back up with a fix. In fairness to Grindr, despite their triaging of security reports needing work, their response after I

managed to get in touch with the right people was exemplary. <u>Here's how they</u> responded when approached by infosec journo Zack Whittaker:

We are grateful for the researcher who identified a vulnerability. The reported issue has been fixed. Thankfully, we believe we addressed the issue before it was exploited by any malicious parties. As part of our commitment to improving the safety and security of our service, we are partnering with a leading security firm to simplify and improve the ability for security researchers to report issues such as these. In addition, we will soon announce a new bug bounty program to provide additional incentives for researchers to assist us in keeping our service secure going forward.

All in all, this was a bad bug with a good outcome: Grindr did well once I got in touch with them, I believe they're making some positive changes around handling security reports and, of course, the bug has been fixed. Oh - and Scott made some new friends \bigcirc .

Epilogue

For the most part this was just another day in security vulnerability land: learn of a bug, let it get fixed, write about it then move onto the next one. But there were 2 things that really set this one apart:

The first thing is the nature of Grindr. I've never used an online dating service before, but I've watched people using the likes of Tinder in the past. Ok, it's full of pretty self-ingratiating photos and some people seem to be a bit pushy but for the most part it appears to me as an outsider to be a pretty representational cross-section of society. Grindr, on the other hand, is extremely overtly sexualised and confrontational. As soon as I loaded Scott's picture in, he started getting requests for "pics". That much I included in the blog post, but what I didn't include is how aggressive things turned. Within minutes of not received requested pics, multiple other users became extremely abusive and

aggressive in ways that shocked even me, and I've seen a lot of shit online over the years! Now keep in mind that Scott's account explicitly said "Test Acc" and we never once engaged in any discussion with anyone, this was simply the result of creating an account then watching what happened and I've gotta be honest here - it was appalling.

The next thing was the response to Scott's pic. Several people contacted both Scott and I privately expressing displeasure at our choice to include that in the post. Their view was that it was somehow a parody of how a gay person might appear and therefore, was not appropriate. I don't even know where to begin with this so maybe just the basics first: the photo was taken by Lars Klint (he wrote the foreword to this book) on the NDC fjord cruise in Oslo. Scott played up to the camera and Lars managed to snap a pic at the perfect time during his most sexy pose. Grindr is a dating site so if we're going to use a pic to illustrate a point, we're going to use the sexiest one available hence Scott's appearance. What frustrates me most about the reactions is the propensity for people to take innocent, good natured and well-intentioned content and turn it into something controversial. Over recent years, this has become an increasingly regular pattern that has honestly sucked a lot of the pleasure out of social media engagements. While I'm on my high-horse, I'll give you an example of this:

In 2019 I went to NDC Minnesota in the US. I did a workshop, the opening keynote and a talk at PubConf, the independently run "variety show and afterparty" that often coincides with NDC events around the world. It's a great night out with 10 speakers doing lightning talks consisting of 5 minutes' worth of slides, each lasting 15 seconds whereupon they auto-progress to the next one. It's at night, at a pub and the organisers make it very clear that it's all off the record and is designed to be edgy and fun. A veteran of many prior PubConfs, I was well prepared and started off early with a slide explaining the nature of the event and that there may be content that offends some people. The slide in question had the following statement on it:

"You know what happens when you get offended? Nothing, now be a fucking

adult and grow up."

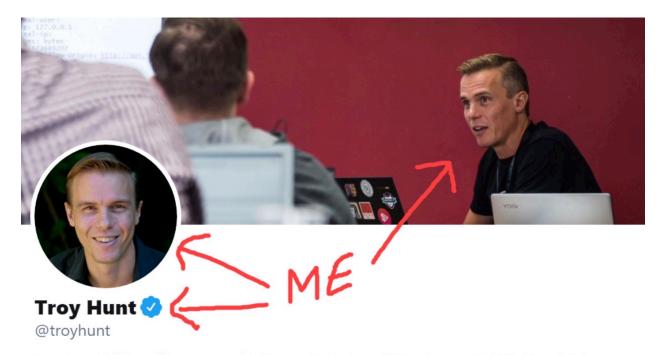
You can guess what happened next: someone got offended at that specific slide. And made a complaint to the organisers. Not even the organisers of PubConf, the independent event running alongside NDC, rather to the NDC organisers themselves. I'm relaying that story again here because although not directly related to the Grindr post, it's the best possible illustration I can give of just how nuts things have gotten when it comes to people being sensitive about the slightest little thing. Whether it be assumptions about Scott's implied sexuality or being offended at being told not to be offended, the world has gone just a bit too crazy.

IF YOU DON'T WANT GUITAR LESSONS, STOP FOLLOWING ME

Rob told me not to include this blog post, something about it throwing a few people under the proverbial bus by including their tweets and ridiculing them. Their public tweets, broadcast to the world and ridiculing me! I can see how some people might pick this up and apply terms such as "punching down", the idea that by virtue of me building a larger audience than those taking shots at me I shouldn't, in turn, criticise them. Bull. Shit. I wrote this blog post because the idea of people dictating what should be on my social media profile is absolutely absurd. That some people then chose to express that view aggressively and in a derogatory fashion on the public timeline makes their inclusion here entirely justified, as far as I'm concerned.

02 NOVEMBER 2020

I 've had this blog post in draft for quite some time now, adding little bits to it as the opportunity presented itself. In essence, it boils down to this: people expressing their displeasure when I post about a topic they're not interested in then deciding to have a whinge that my timeline isn't tailored to their expectation of the things they'd like me to talk about. The key term in that sentence is "my timeline" and as most of this relates to Twitter, there's a very easy way to understand whose timeline you're looking at:



Creator of @haveibeenpwned. Microsoft Regional Director and MVP. Pluralsight author. Online security, technology and "The Cloud". Australian.

1,162 Following **170.2K** Followers

This is me, talking about the things that I find interesting. Ricky Gervais does an amazing job of explaining what I'm about to delve into so do yourself a favour and spend a minute <u>watching this first</u>:



And therein lies the inspiration for the title of this blog. His comedy skit nailed it too: my Twitter timeline is literally just me talking about the things I'm interested in and whilst that might be predominantly technology and infosec stuff, turns out I actually have a life beyond that too. For example, just yesterday I thought it would be nice to take a boat ride and enjoy the impending summer weather down here:





Gold Coast days 😎



3:56 AM \cdot Nov 1, 2020 from Gold Coast, Queensland



Beautiful day out! Lots of lovely responses in the comments too plus, at the time of writing, 144 likes. But not everyone was happy with us being out enjoying the sunshine:





What is this #ShowOff by the privileged tech leaders nowadays (@mitchellh, @troyhunt). We have pandemic and people stuggeling for existence, climate crisis threatening our kids future and we are all about planes, boats and huge houses. Hope I'm not just jeolous or the Twitter Al

9:36 AM · Nov 1, 2020





Read the full conversation on Twitter

In my mind I'm hearing this person in his best Ricky Gervais voice grumbling "but I don't fucking like boats"! Ok, guess you could just ignore them then, would that work? And yes, I know times are tough in many places in the world right now and if that's what you'd like to focus on then by all means, seek out that content. But someone not wanting to see the joy in other people's lives and then berating them for sharing it is just plain stupid.

Don't think this is just a pandemic era phenomenon though; when I bought a new car a few years ago, I was excited and as such, I shared that excitement online:

(Side note: I talked about this particular tweets in <u>my Hack Your Career talk at NDC Oslo</u> a few years ago, deep-linked just to the right spot for your viewing convenience.)

Now, there's one reason and one reason only why I tweeted about the car and I'll summarise it succinctly here:



Troy Hunt @troyhunt · Mar 30, 2017 And this is why we didn't buy a Tesla...







Is there a way to filter that kind of bullsh*t and stick to security/data-breach content _exclusively_? I mean, seriously now...

3:53 AM · Mar 30, 2017



I like cars.

This is not a hard concept to grasp: I post things to my feed I get pleasure from and this person grumbling about "I don't fucking like cars" has absolutely zero impact on my propensity to post more cars in the future (I've posted a lot of car tweets since then). I mean you should see how many pics I post of beer! Oh yeah, apparently that's not on either:

Skimming through the last week of Troy's posts I only see pictures of food, beer, and self promotion

Someone with an audience his size should be using it to help and amplify more important people and issues. Great deal of respect for your work on haveibeenpwned, but disappointed https://t.co/6HdBMYcOnO

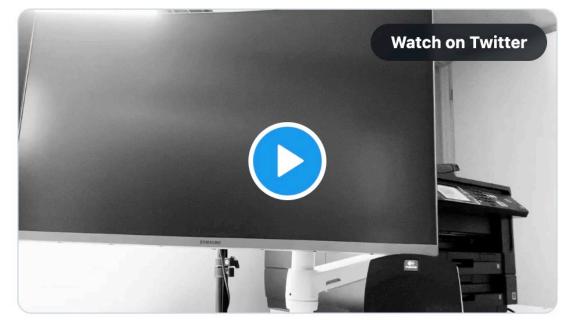
— David McKay (@rawkode) <u>June 5, 2020</u>

I can't recall precisely what the food was but if I felt it was Twitter-worthy, it was probably epic And as for self-promotion, turns out my livelihood does kinda depend on sharing the things I do so that people might take out blog sponsorship or get me to do a talk or allow me to engage in other activities that pay me such that I can buy more food and beer. But David doesn't fucking like food and beer.

I find the sleight against self-promotion in particular a nonsensical position to take on a social media platform I use to amplify my messaging. This aligns very closely to my professional persona as do tweets about how I work or, as I shared quite extensively around the middle of the year, the environment in which I work:









Rich man office, are you showing off? So many IT people in the world could not afford half of one monitor or that ergonomic desk

12:19 PM · Jul 10, 2020



Beyond not so subtly expressing that he doesn't fucking like big monitors, Hakim doesn't really make it clear what can be shown without hurting his feelings. Just one screen? What if it's one of those really slick high-DPI ones that gets really pricey? And what makes that desk "ergonomic"? It's flat on the top and has four legs, is that it? Never mind the fact it's 11 years old and worth nothing and besides, while we're talking about fancy devices:





Replying to @H4k1p

So many people in the world could not afford the pocketsized supercomputer you tweeted that from, but that doesn't seem to bother you

1:04 AM · Jul 23, 2020



It does make me chuckle just a little to see all the likes on that tweet \bigcirc .



It's a constant frustration to see people behave in this fashion, where they pick something that I found interesting, put on it my timeline and because it's not appropriately curated to their personal desires, they sit down and have an angry keyboard rant. Thing is even when I'm bang on topic in terms of the content people expect from me - bang "on brand" as you'll see in a moment - people still get cranky:



Troy Hunt 🔮 @troyhunt · Jun 1, 2020



I'm seeing a bunch of tweets along the lines of "Anonymous leaked the email addresses and passwords of the Minneapolis police" with links and screen caps of pastes as "evidence". This is almost certainly fake for several reasons:



Jennifer Wadella @likeOMGitsFEDAY

Dude, come on. I get it's your brand, but THIS is what you choose to address? Black men are being murdered, but whatever, let's just talk fucking security shit. You want to draw attention to falsehoods help us, point out white nationalists being the perpetrators behind looting

3:45 AM · Jun 1, 2020



Yeah, she pretty much nailed it in terms of being "on brand" because investigating data breaches and writing about their aftermath is pretty much what I've carved out a name for myself doing! But Jennifer doesn't fucking care about disinformation campaigns stemming from data breaches designed to influence public sentiment, and she damn well wants me to know that.

If what I tweet doesn't resonate with you, unfollow me. But don't for a moment think that jumping on the keyboard and telling me you didn't come to my timeline to read what I've put on my timeline is going to influence me one little bit. Right, glad I got that off my chest, I know exactly what I need right now:





Ah, the perfect accompaniment with which to finish this next blog post iii



9:04 AM · Nov 2, 2020 from Gold Coast, Queensland



Epilogue

I've gotten a lot of mileage out of this blog post. Just for kicks, try a Twitter search for the URL and see some of the tweets I've replied to with it. What I find fascinating about the tone of so many of those tweets (and especially the tone

of the ones embedded in the blog post), is that people would never approach me in person and behave that way. That's something I find fascinating about social media communications in general, where so many people get behind the keyboard and behave in a way they know damn well they'd never emulate face to face. I feel like it's a lesson I'd give my children: "don't say anything about other people you wouldn't be prepared to say to their face".

The "guitar lessons" thing has become a bit of an in-joke, with people now often joining my live-streams and complaining that I didn't provide any guitar lessons! I think for me this just became one of those posts that helped me articulate how I wanted to conduct myself online; I'm perfectly happy presenting a more holistic "me" that has all sorts of varied interests and I've no desire to be "the infosec channel", or some myopic interpretation of a single subject. By writing this and then subsequently including it here, I also feel that to use the term from one of those embedded tweets, it's "my brand" in that I'm perfectly comfortable being transparent about how I feel about a subject and that not everyone will agree. Imagine how boring my content would need to be in order to not upset anyone!

PWNED PASSWORDS, OPEN SOURCE IN THE .NET FOUNDATION AND WORKING WITH THE FBI

This is just such a cool blog post, I'm so happy about this one! I'd flagged open sourcing parts of HIBP all the way back in August the previous year and this was finally the time to deliver on the promise. I didn't know how I was going to do it back then and I knew it was going to make more work for myself in the short term, but it was just the right thing to do. You can read about why the .NET Foundation made sense in the blog post and further, the role of the FBI in all this. I think that's what I'm most proud of - having the world's most recognisable law enforcement agency up there alongside the word "pwned" and them playing an active part in sending me compromised data. Mind. Blown

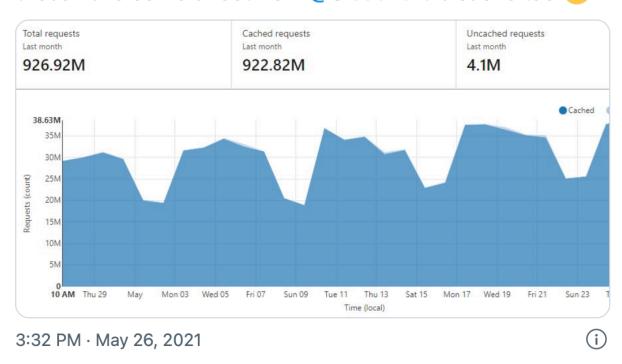
28 MAY 2021

I 've got 2 massive things to announce today that have been a long time in the works and by pure coincidence, have aligned such that I can share them together here today. One you would have been waiting for and one totally out of left field. Both these announcements are being made at a time where Pwned Passwords is seeing unprecedented growth:





Getting closer and closer to the 1B requests a month mark for @haveibeenpwned's Pwned Passwords. 99.6% of those have come direct from @Cloudflare's cache too



That's significant because the sheer volume of requests greatly amplifies the effectiveness of the announcements below. So, keeping in mind this will all be leveraged nearly 1 billion times a month (and much more in the future), read on...

Pwned Passwords is Now Open Source via the .NET Foundation

Back in August <u>I announced that I planned to open source the HIBP code base</u>. I knew it wouldn't be easy, but I also knew it was the right thing to do for the

longevity of the project. What I didn't know is how non-trivial it would be for all sorts of reasons you can imagine and a whole heap of others that aren't immediately obvious. One of the key reasons is that there's a heap of effort involved in picking something up that's run as a one-person pet project for years and moving it into the public domain. I had no idea how to manage an open source project, establish the licencing model, coordinate where the community invests effort, take contributions, redesign the release process and all sorts of other things I'm sure I haven't even thought of yet. This is where the .NET Foundation comes in.

After announcing the intention to go open source, my friend and executive director of the foundation <u>Claire Novotny</u> reached out and offered support, thus beginning a new conversation. I've known Claire for years previously as another Microsoft Regional Director and subsequently as a Microsoft employee and Project Manager on the .NET team. But the .NET Foundation isn't part of Microsoft, rather it's an independent 501(c) non-profit organisation:

The .NET Foundation is an independent, non-profit organisation established to support an innovative, commercially friendly, open-source ecosystem around the .NET platform.

There's a whole page dedicated to the advantages of leaning on the .NET Foundation but in short, they have the answers to all the questions I have no idea about and the dependency HIBP has on the Microsoft stack makes it a natural fit. That it's staffed by a bunch of people I've known and respected for many years and in turn, people that are already familiar with HIBP, makes it a natural fit.

Speaking of natural fits, Pwned Passwords is perfect for this model and that's why we're starting here. There are a number of reasons for this:

- 1. It's a very simple code base consisting of Azure Storage, a single Azure Function and a Cloudflare worker.
- 2. It has its own domain, Cloudflare account and Azure services so can easily

be picked up and open sourced independently to the rest of HIBP.

- 3. It's entirely non-commercial without any API costs or Enterprise services like other parts of HIBP (I want community efforts to remain in the community).
- 4. The data that drives Pwned Passwords is already freely available in the public domain via the downloadable hash sets.

So, I can proverbially "lift and shift" Pwned Passwords into open source land in a pretty straightforward fashion which makes it the obvious place to start. It's also great timing because as I said earlier, it's now an important part of many online services and this move ensures that *anybody* can run their own Pwned Passwords instance if they so choose. My hope is that this encourages greater adoption of the service both due to the transparency that opening the code base brings with it and the confidence that people can always "roll their own" if they choose. Maybe they don't want the hosted API dependency, maybe they just want a fallback position should I ever meet an early demise in an unfortunate jet ski accident. This gives people choices.

That's the open sourcing covered, but what Pwned Passwords *really* needs to be successful is fresh passwords as they become compromised, and this is where the FBI comes in.

The FBI's Feed of Pwned Passwords

As you can imagine, the FBI is involved in all manner of digital investigations. For example, they recently made headlines for their role in taking down the Emotet botnet in conjunction with their law enforcement counterparts in other parts of the world. They play integral roles in combatting everything from ransomware to child abuse to terrorism and in the course of their investigations, they regularly come across compromised passwords. Often, these passwords are

being used by criminal enterprises to exploit the online assets of the people who created them. Wouldn't it be great if we could do something meaningful to combat that?

And so, the FBI reached out and we began a discussion about what it might look like to provide them with an avenue to feed compromised passwords into HIBP and surface them via the Pwned Passwords feature. Their goal here is perfectly aligned with mine and, I dare say, with the goals of most people reading this: to protect people from account takeovers by proactively warning them when their password has been compromised. Feeding these passwords into HIBP gives the FBI the opportunity to do this almost 1 billion times every month. It's good leverage $\ensuremath{\mathfrak{C}}$

I asked the folks there if they'd like to add anything to this blog post and they provided the following statement:

We are excited to be partnering with HIBP on this important project to protect victims of online credential theft. It is another example of how important public/private partnerships are in the fight against cybercrime.

- Bryan A. Vorndran, Assistant Director, Cyber Division, FBI

The passwords will be provided in SHA-1 and NTLM hash pairs which aligns perfectly to the current storage constructs in Pwned Passwords (I don't need them in plain text). They'll be fed into the system as they're made available by the bureau and obviously that's both a cadence and a volume which will fluctuate depending on the nature of the investigations they're involved in. The important thing is to ensure there's an ingestion route by which the data can flow into HIBP and be made available to consumers as fast as possible in order to maximise the value it presents. To do that, we're going to need to write some code. That's right, we're going to need to write some code and thus begins the first piece of open source work for HIBP.

Help Me Build the Code for Password Ingestion

This is a great little first project to distribute to the community and I'm really excited not just about collaboratively working on the code, but that we're doing it in conjunction with a major law enforcement agency to make a positive difference to the world via a free community service. It's wins all round. Here's what I'm thinking:

- 1. There's an authenticated endpoint that'll receive SHA-1 and NTLM hash pairs of passwords. The hash pair will also be accompanied by a prevalence indicating how many times it has been seen in the corpus that led to its disclosure. As indicated earlier, volumes will inevitably fluctuate and I've no idea what they'll look like, especially over the longer term.
- 2. Upon receipt of the passwords, the SHA-1 hashes need to be extracted into the existing Azure Blob Storage construct. This is nothing more than 16 ^ 5 different text files (because each SHA-1 hash is queried by a 5 character prefix), each containing the 35 byte SHA-1 hash suffix of each password previously seen and the number of times it's been seen.
- 3. "Extracted into" means either adding a new SHA-1 hash and its prevalence or updating the prevalence where the hash has been seen before.
- 4. Both the SHA-1 and NTLM hashes must be added to a downloadable corpus of data for use offline and as per the previous point, this will mean creating some new entries and updating the counts on existing entries. Due to the potential frequency of new passwords and the size of the downloadable corpuses (up to 12.5GB zipped at present), my thinking is to make this a monthly process.
- 5. After either the file in blob storage or the entire downloadable corpus is modified, the corresponding Cloudflare cache item must be invalidated. This is going to impact the cache hit ratio which then impacts performance

and the cost of the services on the origin at Azure. We may need to limit the impact of this by defining a rate at which cache invalidation can occur (i.e. not more than once per day for any given cache item).

It's also my hope that the scope of this facility may expand in the future should other law enforcement agencies or organisations that come across compromised passwords wish to contribute. This is just a starting point and I'm really excited to see what direction the community will drive this in.

Next Steps

If I'm completely honest, I don't have all the answers on how things will proceed from here so let me just start with the basics: there's a Have I Been Pwned organisation in GitHub that has the following 2 repositories:

1. Azure Function

2. Cloudflare Worker

The .NET Foundation folks have helped me out with the former and the Cloudflare folks with the latter. They'll continue to help supporting as community contributions come in and as the project evolves to achieve the objectives above re supporting the FBI with their goals. Running an open source project is all new for me and I'm enormously appreciative of the contributions already made by those mentioned above. Bear with me as I navigate my own way through this process and a massive thanks in advance for all those who decide to contribute and support this initiative in the future.

Just one more thing - there's a third repository in that organisation. Because there was so much enthusiasm over this 3D print earlier in the week, <u>I've</u> dropped the .stl into the 3D Prints repository so you can go and grab it and print it yourself. And if you don't have a 3D printer, I'll be sending a bunch of these out I've printed myself to people that make significant contributions to the

project



Comments

10:40 PM · May 24, 2021

Awesome news!

I'm a bit surprised to see that you used blob storage to store password hashes. That means that when you need to add a hash to a prefix file, you need to rewrite the whole file. Wouldn't it be easier to use table storage? Probably with hash prefix as the partition key, and hash suffix as the row key.

_

Troy: I actually started out with Table Storage then moved to Blobs because it was *way* faster for the usage pattern. Full blog post on that here: https://www.troyhunt.com/i-...

—

Probably a good idea is to use table storage and then periodically dump the data into the files that the function uses. So the ingest api uses the table storage for writes and the cloud function uses blob storage for reads

__

Troy: That's a really good point, that could possibly be the most efficient path forward.

Then again, with such a high cache hit ratio I'm wondering how much it really matters if 0.4% of requests are 70ms slower?

On the other hand, I also need to dump all this to a downloadable format at some time and hitting Table Storage in anger does slow down performance on a storage construct I want super fast reads from.

I wonder if the best model might be pretty much what you've said: ingest into Table Storage (either insert or update both SHA-1 and NTLM values with a prevalence count next to them), dump changed SHA-1 prefixes to Blob storage periodically (maybe once a day), flush cache at Cloudflare then once a month, just zip up all the blobs and make them downloadable. Hmm...

Epilogue

After this post I very quickly realised what I'd feared ever since deciding to open source parts of HIBP; it's a lot of work. I was getting heaps of pull requests which was great, but they required my input to review and accept. It was PRs for all sorts of things too; code optimisations, formatting, conventions, project structure and, of course, the new features required by the FBI. Suddenly it was like the project that had begun to cause me so much stress 2 and a half years earlier was now becoming stressful again due to taking actions to delegate work to reduce my stress!

Shortly after this post I announced that Stefán Jökull Sigurðarson would be playing a much more active role in the project, approving PRs and helping guide the development of the ingestion pipeline. I'd known Stefán for years, drank beers with him in various spots around the world and held his work in high regard, especially the efforts he'd made to integrate Pwned Passwords into EVE Online's platform. It wasn't just having someone else around that really helped me out, it was having someone actively involved in coding at a time where I'd been doing less and less of it.

This feels like a nice place to wrap up; my little pet project that began as nothing more than a fun coding exercise made solely to be a free community service is now gradually returning to the community, driven by their generosity and a will to see it sustained for the betterment of everyone. That makes me very happy, and very content $\[\]$

CONCLUSION

e can all break our lives into phases. Sometimes they're defined by significant events in our childhood; I can clearly delineate a very rural life in Australia until I was nearly 14 to my life in the Netherlands for the next couple of years to my arrival in Singapore at 16. Then there were the various relationships, jobs, children, life-changing events and so on. There are more phases to come, and I don't yet know what they are, which is fine. Maybe there's the "Troy being known as the IoT guy" phase. Feasible. Perhaps I become the "person who specialises in [insert JavaScript library here] guy". Unlikely. I don't know, but that's also what keeps it fun.

The current phase of my life is defined by something very different to before. Someone very different to before. Someone who has guided me through the worst of times and in doing so, has given me the best of times. Charlotte. I haven't spoken much about her publicly, and I want to make up for that now.

Charlotte

I mentioned in the epilogue of the post about NDC that I'd met Charlotte for the first time in June 2014. I had very young kids and a stable relationship, she was in a brand new one herself. As beautiful and engaging and frankly, as admired by everyone for all those qualities as she was, 5 years would pass before there was a romance. 5 years of NDCs, speaker dinners, cruises, and then finally, relationship breakdowns. Hers with the guy who also happened to be the NDC photographer (who, incidentally, took many of the photos of me that adorn the pages of my blog), mine with the mother of my children. There wasn't a single

catalyst in either of our breakups, but as I suspect is the case in most relationship endings, a series of small cuts over a prolonged period. Despite living on opposite sides of the world, we came together.

Niall Merrigan was one of the first people I talked to about the relationship. We were sitting in a bar in Oslo bang on 5 years after I'd been pouring my heart out to him in another bar about how shit my job was. He was already very close to Charlotte and couldn't have been happier for me, especially having been witness to the breakdown of both our prior relationships. I mention him here because of what he said immediately after I told him: "Charlotte is always the most beautiful woman in the room who's also the most approachable". I needed that. I felt like I'd won the lottery!



An NDC speaker dinner with Charlotte and Niall (behind her) in 2016

None of this was easy for either of us. Her relationship breakdown was comparably straight forward, not involving kids or marriage despite a tenure of many years. Mine was... opposite. Unsurprisingly, there are all sorts of tricky factors involved when someone new slots into an incumbent life. Someone else's home, their friends, family and kids. Talking to a friend of mine a while ago who'd moved into the home her husband had shared with a previous wife, she confided in me about how it was often the little things she found hard: "It was his ex-wife's handwriting on the spice jars", she said. For Charlotte it was everything from the choice of furniture to which cupboards things went in. Not necessarily major stuff, but reminders that she was coming after someone else whose shadow hadn't entirely departed.

COVID helped. It feels selfish to say that a deadly global pandemic was a blessing, but for us it was. In Feb 2020, we were back in the house I loved so much, right as the virus was starting to change the world. I suddenly went from travelling 243 days in 2019 to going absolutely nowhere. We spent huge amounts of time outdoors in the Aussie sunshine. We could plan for the kids to come and go with absolute confidence there'd be no pressure on me to be away. We met a lot of new neighbours and built out new friendships with people that had no history of our "old" lives. We established stability and consistency the likes neither of us had ever had before, at least not since both our early teenage years.

It took until September 2020 before I was ready to talk publicly about our relationship. It was a difficult thing for me to come to terms with at first because it was a massive change after a 20-year relationship and then because it just took time to feel, well, "normal". That was hard for Charlotte, to be alongside someone who shares so much publicly but had made a conscious effort not to share her role in my life. Along with all the other shit she had to deal with, it's a testament to her that she stuck it out, especially at a time where she had no physical access to her family and friends on the other side of the world and absolutely everything centred around my life here.

But what wasn't difficult was the kids, who from day 1, adored her. She'd sew and go horse riding with Elle, play squash with Ari and have both the most trivial and most insightful discussions with him. She immediately felt like a natural part of my family that was just... right. Not a replacement in any way (neither of us wanted that), but as someone who brought so much more richness to the kids' lives on top of everything they had already. Perhaps what I loved most about Charlotte's role with the kids is that she's so much of what I aspire for them to be; positive, outgoing, fit and active, someone who takes pride in herself and her environment and above all, a genuinely kind soul. She enriched us all.



It's been a huge adjustment for all of us to establish norms in our own "blended" family. In December 2020, we drove a couple of thousand kilometres up the coast of Australia, stopping at many small towns along the way. At

one tiny place that consisted of a convenience store among the gum tress and little else, we all went inside for some cool drinks at the height of Aussie summer. The proprietor was a kindly old lady who looked like she'd spent the better part of 80 years in the local sun, and she greeted the kids as we walked in. "How'd you kids like some lollies? You'll have to ask your mum." Still a new term for her, Charlotte quickly responded with "Oh I'm just their stepmum". Suddenly, the kindly old lady fired up: "You're not *just* their step-mum, you're so much more than that darling." And she was right, but until then I just don't think we recognised the significance of her role in their lives. Ever since that day, there's never been another "just".

We got engaged on New Year's Day, 2021. It was a hard decision for me to make (yes, I proposed), not because of any doubt about my feelings for her, but because of the other "noise" in my life from the previous relationship. I had to somehow shut all that out and not let it cause me to act too fast... or too slow. I didn't want it to influence me in any way whatsoever; that wouldn't be fair to Charlotte, and I wouldn't have been honest with myself. Ultimately, it boiled down to one simple question: if I turned off all the other noise, what would I do? It was an easy answer.



As I write this, it's been 2 and a half years since we've spent more than a few hours apart. As the engagement pic clearly shows, we've travelled, but it's all been domestic and we've done it together. We've been all the way up and down the east coast of Australia, into the middle, down to the bottom, up to Darwin and Kakadu in "The Top End" and there are many more places yet to come. Soon (we all hope), it'll be Norway again and particularly with the kids at an age where they can really absorb culture, I'd like to spend a bunch more time back in Asia. We'll travel a lot more, but never like it was before.

There is a calmness about Charlotte that prevails no matter how weird life gets. She's always balanced, level-headed, and undeterred in her commitment not just to me, but to the kids too. She keeps us sane. She keeps us happy, She makes our family whole \bigcirc .



Epilogue: Charlotte Lyng

When Rob and I talked about writing about Troy I quipped "it's so American to gush over someone... I don't know how to do that". I remember Rob laughing and suggesting I start the chapter just like that so here we go. What can I say? born and raised in Norway we tend to be... a little more *direct*.

"But seriously I have no idea what to write. He's Troy Hunt. He's amazing. I love him... what else...?" Troy's face was in his palm when I said that, and Rob laughed even louder. Then Rob turned the tables on me and asked me in a very

Norwegian way: "Yeah but WHY are YOU with him? He's a lovely guy, sure, but you've stepped into a role as a step-mum for his kids, you moved to the other side of the world, left your life in Oslo behind and your position at NDC to be with him. You've been with him through some very trying times... I want to know why."

I answered the only way I could think of: "It's simple, when you love someone, you do whatever it takes to make it work. I mean... he's my soulmate... how do you describe why you love your soulmate without it sounding *American* about it?"

One thing I've learned whilst writing this is that Troy is extremely good at articulating himself through writing. Me, not so much. But I'm going to try, relating a few stories that, I think, illustrate why I love this man so much.

Sticks and Snakes

I hate snakes and Troy knows that, and he finds it very amusing how jumpy I am when we're in an area where there might be some which, it turns out, is *everywhere* in Australia.

I blame my sister for this by the way. Growing up she would hide under the stairs at home waiting for me to come downstairs so she could jump out and scare me and it worked every time.

I'm not a naturally fearful person, but if you hide under the stairs and jump out at me – I'll probably kick you right after I scream. I was just conditioned that way by my merciless older sister.

The other thing you need to understand is that, as a Norwegian, I've also been conditioned to believe that Australia will kill me. *Nyheter*, or "the news" in Norway, always seemed to feature the latest shark attack or hail the discovery of a drop bear. Troy assures me that's not a thing, but sometimes I can convince myself one is hiding under our stairs waiting for me.

Troy finds all of this extremely amusing.

I think he's right – it *is* amusing. But hiking with him still requires courage on my part, and a desire to see Troy jump for once. Two years ago, Troy and I went on a hike to Barrenjoey Lighthouse in Palm Beach, Troy grinned and said: "It's snake season so look carefully when you walk through the bushy areas. I'm serious, some of the deadliest snakes in Australia can be found here and they like to come out on the trail when the sun is out".

Wonderful. Are there drop bears out here too? I didn't entirely believe him, and my courage just seemed to fade on its own as I started to walk faster. Fast enough to put some distance between Troy and me. Which only made things worse.

I was constantly looking down, scanning the ground for slithery things. I'm not sure what I would have done if I saw one – I mean by then it would probably be far too late, and I'd be dinner! Between my constant scanning and my increased pace, I lost track of Troy completely.

Not a good thing. I found out two things that day: I'm still very easy to scare and it's also very easy for Troy to become a mischievous 15-year-old – not too far off from the picture on the front of this book.

From 20 or so metres behind me, Troy had picked up a stick and he threw it next to my right foot, in a bush that was especially dry. The sound the stick made as it struck the dry branches was incredibly snake-like and in an instant, I could see a 12-metre black mamba deciding which leg would be easiest to drive its fangs into.

I ran, *fast*. No – that's not true – I shrieked and ran. I was sure Troy could take care of himself – he's Australian and Australians never seem to be victims of these way-too-deadly creatures. Probably best to be sure, so I turned around and felt a wave of relief and a small wave of *sinne* - "anger" which quickly faded. Troy was standing there laughing, camera in his hand, filming the entire thing.

I initially wanted to smack him, but I couldn't help myself from laughing too. The look on his face, the sound of his cackling... I don't know how he does it, but he has this ability to change any setting into a fun one, even at the worst of times. He really is that boy on the cover of this book – bright-eyed, curious, mischievous, loving and kind. Both of his kids (Ari and Elle) really love that about him too. I often joke that I live with three children, and I wouldn't have it any other way!

A Random Coincidence

We were on our way to the speaker's event for NDC Oslo (a cruise of the Oslo Fjord) in 2015 and Troy and I were chatting aimlessly about life. To be honest I don't remember what the conversation was about; It might have been work (I was one of the organisers of NDC Oslo so chatting up the speakers was part of my job) or the weather... I really don't know. What I do know is that talking with Troy was easy then and it's still easy now. It doesn't have to be sensational conversation either; just two people enjoying each other's company.

Right then, one of the speakers (Adam Cogan) asked if he could take a picture of us together.



It's a cliché to say" if I only knew then what I know now" but in every sense, that's what I think when I see this picture of us together. A distant American friend remarked that we looked like "a beautiful Norwegian couple!" when Adam tagged us in this picture on Facebook in 2015 to which I replied "ha-ha, – we are not a couple! He is one of the speakers at the conference I'm organising".

We were professional acquaintances at best, although I suppose I thought of Troy as a friend too. There was nothing romantic between us at all – yet when I look at this picture, I can't help but reflecting upon the Troy I knew then, versus the Troy I know today. Having dealt with many different *speaker personalities* through my work, Troy is very much *what you see is what you get* kind of person, he doesn't have an "online" and an "offline" persona like a lot of others do. He is the same amiable person regardless of which setting you meet him in. He is one of the most inclusive and generous people I know, always there with a smile.

However, one thing that has become more evident over time, is how hard he pushes himself. He's faced several hardships during the course of our relationship, but he never seems to get bogged down in self-pity, instead he throws himself into work. He is passionate about what he does, he sets the bar high, and absolutely excels at what he does. From the start of our friendship up to this day, I've always admired Troy for who he is, his accomplishments, and I genuinely believe the world is a better place for having him in it.

As I write this, our wedding day is almost here. The ring on my finger promises a lot of things, suggests a life that's to come. Troy started this book wondering about himself at 18, so I'll end it by wondering about the two of us many, many years in the future, looking back at our time together, reading this book maybe, and possibly saying "if we only knew then what we know now".

Our journey together will be filled with many easy days together and, hopefully, no snakes or drop bears but, if there are, I'm sure Troy will be there, laughing about it.

Norwegians have a few different ways of expressing love, unlike Australians or Americans who might say "I love you" to a friend, parent, or partner. Being exacting Scandinavians, we have a specific phrase that we use sparingly – in fact I would say *rarely* because it's that precious. I warned Rob that I wouldn't gush about Troy, and I'm going to hold to that. But telling someone you truly, deeply love them is not gushing – it's true.

Troy: jeg elsker deg, always.