Security Now! #693 - 12-11-18

“Internal Bug Discovery”

This week on Security Now!

This week we take a look at Australia's recently passed anti-encryption legislation, details of a couple more mega-breaches including a bit of Marriott follow-up, a welcome call for legislation from Microsoft, a new twist on online advertising click fraud, the DHS is interested in deanonymizing cryptocurrencies beyond Bitcoin, the changing landscape of TOR funding, an entirely foreseeable disaster with a new Internet IoT-oriented protocol, a bit of errata and some closing the loop feedback from our truly terrific listeners... and then we look at a case where a prominent company discovered one of their own bugs and acted responsibly -- again -- and what that suggests for everyone else.

This week’s Security Now! Picture of the Week
Needed a page all to itself, so…

“Whoa!! Apparently “LogMeIn” has a backdoor!!”...


LogMeIn
USE MAIN
ENTRANCE
EXIT ONLY

Security News

Australia's Telecommunications Assistance and Access Bill 2018
Last Thursday, Australia's House of Representatives has finally passed the "Telecommunications Assistance and Access Bill 2018," less formally known as the Anti-Encryption Bill. Once Australia's upper house votes the bill into law, which =IS= expected, since the bill enjoyed wide bipartisan support, it would go into effect during the next session of Australia's parliament in early 2019.

The Bill would provide law enforcement the legal means to compel Google, Facebook, WhatsApp, Signal, and any others to assist them in accessing the encrypted communications of their products and platforms. The Bill adds new provisions to Australian law, specified in three so-called "notices":

- TAR - Technical Assistance Request - is a "notice" to request tech companies for providing "voluntary assistance" to law enforcement, which includes "removing electronic protection, providing technical information, installing software, putting information in a particular format and facilitating access to devices or services."

- TAN - Technical Assistance Notice - is a notice requiring rather than requeting technology companies to provide assistance they are already capable of providing that is reasonable, proportionate, practical and technically feasible, giving Australian agencies the flexibility to seek decryption of encrypted communications in circumstances where companies have existing means to do it (like at points where messages are not end-to-end encrypted).

- TCN - Technical Capability Notice - is issued by the Attorney-General requiring companies to "build a new capability" to decrypt communications for Australian law enforcement.

Collectively, these so-called "notices" would compel tech companies to modify their software and service infrastructure to backdoor encrypted communications and data that could otherwise not be obtained.

https://parlinfo.aph.gov.au/parlInfo/download/legislation/billsdgs/6370016/upload_binary/6370016.pdf

Following earlier industry consultations, the Government released an Exposure Draft of the Bill on 14 August 2018 and sought public submissions by 10 September 2018. The Department of Home Affairs (DoHA) received almost 16,000 submissions, of which over 15,000 were classified as standard campaign responses, 743 were ‘unique individual responses classified as appropriate for consideration’ and 55 were ‘considered substantive submissions from industry groups, civil society, government bodies and individuals’. While some stakeholders raised concerns about other schedules, the majority of submissions focused primarily or exclusively on Schedule 1 of the Exposure Draft (industry assistance).

Many stakeholders provided submissions that included general and specific recommendations on the proposed industry assistance scheme, including the IGIS, AHRC, LCA and applied cryptography academics Chris Culnane and Vanessa Teague. There was significant concern that

the scheme in its current form has a very wide application and that amendments to offer greater definition, narrow the scope or clarify processes are necessary.

From a technology perspective, Apple submitted that Schedule 1 ‘remains dangerously ambiguous with respect to encryption and security’. Further, Apple stated: We encourage the government to stand by their stated intention not to weaken encryption or compel providers to build systemic weaknesses into their products. Due to the breadth and vagueness of the Bill’s authorities, coupled with ill-defined restrictions, that commitment is not currently being met. For instance, the Bill could allow the government to order the makers of smart homespeakers to install persistent eavesdropping capabilities into a person’s home, require a provider to monitor health data of its customers for indications of drug use, or require the development of tool that can unlock a particular user’s device regardless of whether such [a] tool could be used to unlock every other user’s device as well… While we share the goal of protecting the public and communities, we believe more work needs to be done on the Bill to iron out the ambiguities on encryption and security to ensure that Australian are protected to the greatest extent possible in the digital world.

For its part, the Australian government argues the new legislation is important for national security and an essential tool to help law enforcement and security agencies fight serious offenses such as crime, terrorist attacks, drug trafficking, smuggling, and sexual exploitation of children.

Since the bill had support from both major parties (the Coalition and Labor), the upper house could vote in support of the Assistance and Access Bill to make it law, which is expected to come into effect immediately during the next session of parliament in early 2019.

The Bill states that the tech companies can't be compelled to introduce a "systemic weakness" or "systemic backdoor" into their legit software or hardware, or "remove electronic protection," like encryption to satisfy government demands.

Instead, the new legislation contains measures aimed at facilitating lawful access to information through two avenues—"decryption of encrypted technologies and access to communications and data at points where they are not encrypted."

The Bill stipulates: "We encourage the government to stand by their stated intention not to weaken encryption or compel providers to build systemic weaknesses into their products." So without forcing companies to break the encryption in their software, Australian law enforcement is looking for ways to snoop on messages before they are encrypted, or to read them once they're decrypted on the other users' end.

This would require assistance from providers of the software and services, including Apple, Samsung, Google, WhatsApp, Signal, iMessage, and Telegram, though it remains to be seen whether and how tech companies cooperate with the new Australian laws.

Five Eyes Nations: Responses to "Going Dark"

Since Australia is a member of the Five Eyes alliance along with the United States, United Kingdom, Canada, and New Zealand, which last month declared that "privacy is not an absolute" and the use of end-to-end encryption "should be rare," the new bill could be a stepping stone towards new encryption laws in other nations as well.

The Bill also claims that without the new legislation, law enforcement agencies face the problem of "going dark"—which is, as we know, a term used by the FBI and U.S. Department of Justice (DoJ) to describe the condition of the increasing use of encrypted communications.

Australian Prime Minister Malcolm Turnbull has previously made his position on encryption clear last year, saying "The laws of mathematics are very commendable, but the only law that applies in Australia is the law of Australia."

Apple: Encryption is Simply Math

Apple responded to the new bill by making a submission to the Australian government month ago, saying "Encryption is simply math. Any process that weakens the mathematical models that protect user data for anyone will by extension weaken the protections for everyone."

Apple wrote: "It would be wrong to weaken security for millions of law-abiding customers in order to investigate the very few who pose a threat."

This week in Mega-Breaches...

The Attorneys general of the 12 US states: Arizona, Arkansas, Florida, Indiana, Iowa, Kansas, Kentucky, Louisiana, Minnesota, Nebraska, North Carolina, and Wisconsin have gathered to file a lawsuit against Indiana-based Medical Informatics Engineering (MIE) and its subsidiary NoMoreClipboard (NMC) this week. The states, each of which have residents affected by the breach, are negotiating a payout with the company.

MIE sells web-based electronic health record services to healthcare providers via NoMoreClipboard's Webchart web-based portal.

Way back on May 7th of 2015, hackers stole the very personal information of 3.9 million people from MIE’s back-end database using a simple SQL injection attack. Not only were names, addresses and social security numbers stolen, but also health data including lab test results, health insurance policy information, diagnoses, disability codes, doctors’ names, medical conditions and the names and birth statistics of children. In other words... everything.

The joint complaint accuses MIE of failing to properly secure its computer systems, not telling people about its system weaknesses, and then failing to provide timely notifications of the incident. MIE never bothered to actually encrypt sensitive information, even though it said it did. It also provided public accounts sharing the passwords “tester” and “testing” which were established so that a client’s employees didn’t have to log in with a unique user ID. (Yeah, why bother with all that?)

Penetration testers uncovered the issue, and highlighted the risk, but the lawsuit says that MIE took no action. Using one of these test accounts, the thieves explored the health record database with SQL injection attacks to gain further access to more privileged accounts.

MIE allegedly didn't have data exfiltration alarms in place. It was a network performance monitoring alarm that finally raised the red flag because the attackers were dumping so many large records from the database, at such volume, that it choked off the network bandwidth. And the data exfiltration continued while administrators investigated the incident. "Hmmmm... I wonder why the Internet is slow today?"

The States allege that once the breach was discovered, MIE only had a draft incident response plan, and there was no evidence that it followed that in any case. And they added that notifications were inadequate. MIE discovered the breach on 26 May 2015, and informed the public of the breach via a notice on its website (only) fifteen days later, on June 10th. The company then finally began email notifications another five weeks later, on July 7th... and finally sent letters in December.

For their part, the 12 States claim that MIE and NMC violated federal HIPAA legislation protecting the privacy of health information. They also accuse MIE of breaking 27 state-level laws concerning data breach notification, abusive and deceptive practices, and personal information protection. The states are proposing a consent decree to clear up the matter before getting into litigation. This calls for an as-yet undefined payout from MIE, along with its commitment to follow several security measures, including the use of multi-factor authentication, not making generic accounts accessible via the internet, using strong passwords, training staff properly in cybersecurity, using a security incident and event monitoring (SIEM) solution, and putting SQL injection attack detection measures in place.

The company will also have to conduct regular security audits with help from a qualified professional, file reports, and take action on them. In short, the settlement asks the company to do what any competent cybersecurity team charged with protecting sensitive data should be doing.

In their coverage of this, Sophos notes the collaborative nature of the settlement. They wrote: "As voices call for stricter federal privacy protection laws, this could be a sign that states are getting fed up with these mega-breaches and are taking things into their own hands." In October, Uber settled with all 50 states over the handling of its 2016 data breach, paying $148m. Sophos wonders whether this suit might herald more coordination between attorneys general to hold companies accountable?

If so, let's hope that other organizations sit up and take note. In today's security climate, security is either going to be paid for now or later. Now is always less painful... and much less expensive!

Databases? We've got your Databases! Get'em right here!
Bob Diachenko, the director of Cyber Risk Research at Hacken, had been tracking a publicly accessible instance of MongoDB, which could be accessed without authentication, for some time.

The database contained 66,147,856 unique records containing full name, personal or professional email address, user's location details skills, phone number, and employment history. A link to the individual's LinkedIn profile was also present, all of which led him to believe that the records were likely scraped data from LinkedIn profiles.

Bob initially discovered the collection in October, in a repository called "database" that contained 49 million records, which was part of a larger discovery that revealed over 120 million records exposed and another, was managed by a company in Florida with 22 million records that included the email addresses, names, and the area where a candidates sought a job. The other collection had 48 million entries with names, work email addresses, phone number and employee details.

He was unable to determine the owner of the 66+ million record database, but he told Bleeping Computer that it is no longer online. This does not exclude the possibility of popping on the web again, though.

If any of our listeners are concerned, the scraped data has been uploaded to Troy Hunt's HaveIBeenPwned service which will allow users to check whether their personal information has been exposed.

And regarding the legality of web scraping for personal data, Diachenko says that it is legal to copy what is publicly available, but it should not be used against the best interests of the owner, which is considered an offense.

He told Bleeping Computer: "Since the data displayed on websites is meant for public consumption, it is legal to copy the information to a file on your personal computer. However, if that information is used in any way that goes against the best interests of the owner, then it is not legal."

Because there is the risk of the personal data to be used against you, the researcher recommends sharing only the "bare minimum" when creating an online profile or account.

-----> And so this, too, seems to be the natural consequence of having a global network of commercial entities, the incredible dropping cost of long term high volume mass storage, and the inherent resale value of the aggregated personal information of millions of individuals. We're no longer the consumer who is buying... the details of our existence is the product that's being sold... without our knowledge or permission.

I get it that many people will not and do not care about the circulation of their personal information. But for those who do, wherever possible minimizing sharing DOES make good sense.

... Which leads us nicely into our next bit of news about Microsoft and facial recognition...

Microsoft asks for legislated Facial Recognition limits - Now!
https://blogs.microsoft.com/on-the-issues/2018/12/06/facial-recognition-its-time-for-action/

Facial recognition: It's time for action | Dec 6, 2018 | Brad Smith - President

This was preceded by a post which lays out the problem as Microsoft sees it:

https://blogs.microsoft.com/on-the-issues/2018/07/13/facial-recognition-technology-the-need-for-public-regulation-and-corporate-responsibility/

Facial recognition technology: The need for public regulation and corporate responsibility
Jul 13, 2018 | Brad Smith - President

<quote> "Some emerging uses are both positive and potentially even profound. Imagine finding a young missing child by recognizing her as she is being walked down the street. Imagine helping the police to identify a terrorist bent on destruction as he walks into the arena where you're attending a sporting event. Imagine a smartphone camera and app that tells a person who is blind the name of the individual who has just walked into a room to join a meeting.

But other potential applications are more sobering. Imagine a government tracking everywhere you walked over the past month without your permission or knowledge. Imagine a database of everyone who attended a political rally that constitutes the very essence of free speech. Imagine the stores of a shopping mall using facial recognition to share information with each other about each shelf that you browse and product you buy, without asking you first. This has long been the stuff of science fiction and popular movies – like "Minority Report," "Enemy of the State" and even "1984" – but now it's on the verge of becoming possible.

Perhaps as much as any advance, facial recognition raises a critical question: what role do we want this type of technology to play in everyday society?

The need for government regulation

The only effective way to manage the use of technology by a government is for the government proactively to manage this use itself. And if there are concerns about how a technology will be deployed more broadly across society, the only way to regulate this broad use is for the government to do so. This in fact is what we believe is needed today – a government initiative to regulate the proper use of facial recognition technology, informed first by a bipartisan and expert commission.

So what issues should be addressed through government regulation? That's one of the most important initial questions to address. As a starting point, we believe governments should consider the following issues, among others:

- Should law enforcement use of facial recognition be subject to human oversight and controls, including restrictions on the use of unaided facial recognition technology as evidence of an individual's guilt or innocence of a crime?

- Similarly, should we ensure there is civilian oversight and accountability for the use of facial recognition as part of governmental national security technology practices?

- What types of legal measures can prevent use of facial recognition for racial profiling and other violations of rights while still permitting the beneficial uses of the technology?

- Should use of facial recognition by public authorities or others be subject to minimum performance levels on accuracy?

- Should the law require that retailers post visible notice of their use of facial recognition technology in public spaces?

- Should the law require that companies obtain prior consent before collecting individuals' images for facial recognition? If so, in what situations and places should this apply? And what is the appropriate way to ask for and obtain such consent?

- Should we ensure that individuals have the right to know what photos have been collected and stored that have been identified with their names and faces?

- Should we create processes that afford legal rights to individuals who believe they have been misidentified by a facial recognition system?

Android Ad-Click Fraud campaign pretends to be Apple devices.
A long time ago somewhere here or on the TWiT Network we noted the interesting fact that in this tightly optimized online advertising land, ads clicked by iOS device users were considered to be more valuable than those click by Android users, so advertisers shelled out fractionally more moolah for iOS clicks than for clicks from non-iOS devices.

So it should come as little surprise that since there is already no honor among thieves, Sophos recently uncovered 25 Apps on the Google Play store that were spoofing their highly spoofable "User-Agent" headers to lie about the device whose user was supposedly (though not actually) clicking on webpage ads. The whole thing was being done in a hidden off-screen window so the App's users were blissfully unaware.

Marriott will reimburse past guests for the cost to replace compromised passports
A Marriott spokesperson said the hotel chain is working on a way to reimburse some prior guests for the costs of obtaining replacement passports if they can show they've been the victims of fraudulent operations where the passport number was involved.

The statement came after New York Senator Chuck Schumer publicly called on Marriott to pay for people's passport replacement fees. Schumer said: "The data breach at Marriott compromised millions of travelers' U.S. passport info. A new passport costs $110. So Marriott must personally notify customers at greatest risk. And Marriott should pay the costs of a new passport for victims who request it."

Whereupon, Connie Kim, Marriott spokesperson wrote in an email to The Washington Post: "We are setting up a process to work with our guests who believe that they have experienced fraud as a result of their passports being involved in this incident. If, through that process, we determine that fraud has taken place, then the company will reimburse guests for the costs associated with getting a new passport."

As we covered last week, Marriott said hackers had gained access to the hotel guest reservation system used at Marriott subsidiary Starwood, and all of the many related StarwoOD properties worldwide.

The US DHS hopes to de-privacy-ify the explicitly private Zcash and Monero currencies
https://www.fbo.gov/utils/view?id=f0e31ab37561cac3cc4a4ab88d9059b0

According to a pre-solicitation document, the DHS wants to know if this is possible, before filing an official solicitation request later down the line, to obtain the same sort of transaction information for Zcash and Monero that they've been able to obtain for Bitcoin.

As we know, law enforcement has had great success in tracking Bitcoin transactions. Although Bitcoin is cryptographically secured against external tampering, Bitcoin transactions are not private since they are each recorded in a public, though tamper-proof, ledger.

The DHS presolicitation document begins by explaining:

The Department of Homeland Security (DHS) Small Business Innovation Research (SBIR) Program, comprised of the Science and Technology (S&T) Directorate’s SBIR Program and the Countering Weapons of Mass Destruction (CWMD) Office SBIR Program, invites small business concerns to review this pre-solicitation notice, which is intended to lead to the FY19 DHS SBIR Phase I solicitation. This notice is not a solicitation or Request for Proposals. This notice is merely an opportunity for interested parties to comment on or request information about the attached topic areas.

The last item on the list, which includes a bunch of other interesting techie interests, is:

OBJECTIVE: Design a product to support the implementation of block chain based forensics, data analysis, and information sharing.

DESCRIPTION: Blockchain and Distributed Ledger Technology (DLT) are emerging technologies being leveraged for a wide range of commercial and governmental applications. The most well-known use case would likely be Bitcoin, within the newly emerged cryptocurrency arena, which has spurred further interest and developments. Prior efforts have addressed Bitcoin analytics, which covers only a limited scope within the realm of cryptocurrencies. This proposal seeks applications of blockchain forensic analytics for newer cryptocurrencies, such as Zcash and Monero. And, ongoing research within the field also contributes to new technological implementations and techniques that continue to multiply the specific types of consensus, privacy, security, and proof mechanisms.

A key feature underlying these newer blockchain platforms that is frequently emphasized is the capability for anonymity and privacy protection. While these features are desirable, there is similarly a compelling interest in tracing and understanding transactions and actions on the blockchain of an illegal nature. To that end, this proposal calls for solutions that enable law enforcement investigations to perform forensic analysis on blockchain transactions. This analysis can be approached in any number of ways and may consider different data situation use cases depending on whether additional data from off-chain sources are available. Furthermore, with the proliferation of new blockchain variants, the desired solution should either attempt to show generality or extensibility, or at least provide working approaches to treating newer blockchain implementations.

PHASE I: Design a blockchain analysis ecosystem or modify an existing one, that enables forensic analysis for homeland security and law enforcement applications for cryptocurrencies, such as Zcash and Monero. Produce an architecture that shows how system components can be upgraded or interchanged for an extensible and forward-looking solution that can be maintained for use with emerging blockchain networks. Demonstrate or discuss implementation feasibility with respect to: concept of operations, governance, algorithms, costs, and security. Identify risks to privacy, security, and technology and develop risk mitigation strategies.

PHASE II: Prototype and demonstrate the blockchain forensic technologies designed during Phase I. The demonstrations will include three (3) use cases determined by DHS/S&T and will involve the analysis of suspicious transaction without external data, with external data, and on another blockchain platform. A technical report detailing the results and improvements made to enhance the technology will be provided after each demonstration.

PHASE III: COMMERCIAL OR GOVERNMENT APPLICATIONS: The proliferation of blockchain technology beyond the cryptocurrency arena has drawn interest from all other sectors, with new proposed blockchains for everything from banking, charitable donations, supply chain tracking, to automatically executing "smart contracts". These technologies stand to radically transform operations in government and the private sector. Because of the significant impact in areas such as governance, data sharing agreement enforcement, and encrypted analytics interchanges, there are a wide variety of applications in government and the commercial marketplace that can benefit from successful product development. Blockchain forensic analytics for the homeland security enterprise can help the DHS law enforcement and security operations across components as well as state and local law enforcement operations. Private financial institutions can likewise benefit from such capabilities in enforcing “know your customer” and anti-money laundering compliance.

In the case of Bitcoin, tracking its transactions, with the help of the private sector, allowed US law enforcement authorities to determine that 95 percent of all ransomware ransom payments were cashed out and converted into fiat currency through the BTC-e cryptocurrency exchange. An international arrest warrant for the owner of the BTC-e portal was issued. But the US lost the extradition fight with Russia.

The changing funding of Tor

- Our coverage of Tor in the past.
- Internet: Privacy, Security, Anonymity
- Point-to-point connections.

The TOR project is independent. Its technology is well known and understood and solid, and its underlying code base is open source and well vetted. But it's always been a bit of a worry that the original development funding and the vast majority of funding for many subsequent years was from the US State Department. It's difficult to see how the US government -- or any government for that matter -- would want to be actively encouraging and supporting anonymity which, while there are many useful and noble applications, also undoubtedly has a dark underbelly downside.

So it's with some sense of relief that the Tor Project's recently released funding disclosure now shows that US support is down to just over half, at 51% with the balancing 49% now being supplied by the private sector.

The Tor Project's co-founder and original developer of the Tor software, Roger Dingledine, said: "In terms of percentages, while 2015 saw 85% of our funding coming from US government sources, 2016 saw the fraction drop to 76%, and in 2017 we're down to 51%."

The project's Non-profit Tax Exemption Form 990 further reveals that the largest source of non-private government funds weren't donations from government agencies, but were research grants from other government-backed organizations.

https://www.torproject.org/about/findoc/2017-TorProject-Form990.pdf

And even so, Roger noted that "I should take a brief moment to explain how funding proposals work, for those who worry that governments come to us wanting to pay us to do something bad. There is never any point where somebody comes to us and says 'I'll pay you $X to do Y'. The way it works is that we try to find groups with funding for the general area that we want to work on, and then we go to them with a specific plan for what we'd like to do and how much it will cost, and if we're lucky they say ok."

The 2017 break down of funding looks like this:

$798,029 worth of funding from the US government-backed Radio Free Asia
$635,504 from the similarly US government-backed SRI International
$594,408 from the Swedish International Development Cooperation Agency (SIDA).
$548,151 from the US National Science Foundation (NSF)

And the US State Department which, as we noted, financed Tor's initial development and has sustained the project through its first years by covering most of its costs, has been reducing its involvement in Tor Project funding. It donated:

$199,071 in 2015
$218,796 in 2016
$133,061 in 2017

At the same time, last year the Tor Project raised a record $425,709 from its users, which is more than twice the funds it raised from users in 2016. But this year it looking rather bleak so far at roughly $95,000 so far.

Mozilla has been a good partner as well. They have pledged matching funds for user donations and Mozilla has overall stepped up, boosting its contributions from $24,500 in 2016 to $522,188, last year, and will likely be a top contributor in 2018, too.

DuckDuckGo also contributed $25,000 to the Tor Project

And the Project has also received in-kind service donations such as free could computing, free hosting, and volunteer coding, translation and legal services... all totaling $806,372.

All told, the Project's 2014 total revenue was $2.5 million. It grew to $3.3 million in 2015, dropped a tad to $3.2 million in 2017, and hit an all time high of $4.2 million in 2017.

Bleeping Computer reports that despite reporting its highest year of income to date, the Tor Project is still scraping by, mainly because its costs have also gone up as well. However, Tor Project has used its income better than most. This year, the organization released a major redesign of its desktop browser but also launched its first official Android browser version.

Roger Dingledine said: "Tor's budget, even at the 2017 level, remains modest considering the number of people involved and the impact we have. And it is dwarfed by the budgets that our adversaries are spending to make the world a more dangerous and less free place."

Tor's donation's page is: https://donate.torproject.org/

And it's live. If you watch it for a while you might see it bump up a bit.

It's really nice to know that in a network designed NOT to provide anonymity, such a facility exists. It really does provide strong anonymity which, as we've discussed previously, can only be penetrated at great cost and with tremendous resources.

An entirely foreseeable disaster: The UDP-based CoAP protocol.

http://coap.technology/

<quote> CoAP / RFC 7252 -- Constrained Application Protocol

“The Constrained Application Protocol (CoAP) is a specialized web transfer protocol for use with constrained nodes and constrained networks in the Internet of Things. The protocol is designed for machine-to-machine (M2M) applications such as smart energy and building automation.”

Now here's the bad news... It supports the use of unauthenticated UDP packets in a simple Query/Reply mode... and a small query can return a ~50x larger reply.

That's everything we need to know on this podcast, to know that it's a dream come true as the next big DDoS amplification protocol.

CoAP is essentially a compactified HTTP that =CAN= support DTLS-style certificate-based endpoint authentication and encryption. But when that's done it's no longer tiny and simple and lightweight. So all that extra goodness has been eschewed in favor of just sending a query and returning a reply.

CoAP is new, having been ratified into a standard only 4 years ago in 2014. So we would expect uptake to follow the creation of libraries and manufacturing and deployment.

What does Shodan have to say about it?

After several initially quiet years, the number of publicly exposed, accessible and answering CoAP devices has exploded since a little over a year ago in November 2017.

After first appearing on the map a year ago at ~6,500 devices. The next month it was at 26,000. Then, by May of this year that number was 278,000 devices... and today we're in the 580,000-600,000 range, according to Shodan.

Dennis Rand, founder of eCrimeLabs who has been watching this happen believes that the reason for this explosion is CoAP's use as part of QLC Chain (formerly known as QLink), which is a project aimed at building a decentralized blockchain-based mobile network using WiFi nodes throughout China.

China: 573,910
US: 3,831
Russia: 1,675
Canada: 327
Germany 193

And this recent rise in readily available and poorly secured CoAP clients hasn't gone unnoticed. Over the past few weeks, the first DDoS attacks carried out via CoAP have started to appear. ZDNet reported that a security researcher who deals with DDoS attacks, but who couldn't share his name due to employment agreements, stated that CoAP attacks have occurred occasionally during the past few months, with increasing frequency, reaching 55Gbps on average, and with the largest one clocking at 320Gbps.

55Gbps is an order of magnitude larger that the average size of a normal DDoS attack, which is 4.6Gbps, according to DDoS mitigation firm Link11.

Of the 580,000 CoAP devices currently available on Shodan today, the same researcher told ZDNet that roughly 330,000 could be (ab)used to relay and amplify DDoS attacks with an amplification factor of up to 46 times.

And, of course, this was all foreseeable.

One way to resolve this trivially would have been to have the initiator first ask for permission by sending a small permission request. The recipient would encrypt the requesting IP with its own locally unique private key and return it as a permit. Then, only requests containing the proper permit for the requesting IP would be replied to. This would require an extra UDP packet round trip, but only once per endpoint pair, so the overhead would be minimal. And the recipient never needs to retain any state. It just issues permits for permission requests and checks the permits on incoming data requests. So you would still get UDP reflection, but not any amplification. So the protocol would be of no interest to attackers.

Alert (TA14-017A)
UDP-Based Amplification Attacks
https://www.us-cert.gov/ncas/alerts/TA14-017A

Errata

André Domnick @secorvo
Subject: Sennheiser Root CA - Google Chrome Certificate Pinning
Date: 10 Dec 2018 06:45:40
:
Hi Steve,

Thanks for mentioning our small finding, but just a short comment on the Chrome stuff:

It is often mentioned in the media that Google Chrome is performing Certificate Pinning for the Google domains, but Google appently has backed from the idea. In our report you can find a picture of the current Google Chrome accepting the certificate. Additionally we did not detect any unusual behaviour of Chrome when presenting our cert to the browser ;-)

Best Regards
André

Closing The Loop

Greg K
Location: Columbus, OH
Subject: CA Certificates
Date: 08 Dec 2018 10:09:14
:
I'm fine with your rants about adding certs to trusted ca stores, and "middle boxes". However, I'm an advanced user. If I want to use my own internal Root CA or my own "middle box" to knowingly scan my own web traffic. I should be able to do that. I should not be restricted from adding whatever I want to my trusted cert stores on my network through group policy or whatever means. So I agree with your comments, but I'm opposed to restricting the ability to do these things; but I know what I'm doing but others do not. So please be careful and responsible when you call for banning these activities to those individuals or organizations that can do it responsibly. I know you didn't exactly do that, but I would hope that you are being careful with your opinions.

Craig
Location: Edinburgh (no accent, please, Leo)
Subject: Self-signed cert surprise all the way from Azeroth
Date: 07 Dec 2018 11:00:20
:
Given your audience, I bet I'm not alone in alerting you to this, but thought it was worth flagging up. After you mentioned SigCheck I ran it, expecting to be clean as I don't install much, and was surprised to find one listed from Blizzard Battle.net.

Yes, I am a long-time Warcraft addict. I could quit any time I like. Honest.

I wondered what it was doing there and came across this post from a Blizzard employee:
https://us.battle.net/forums/en/bnet/topic/20760626838

I don't think it's anything to really worry about, as the last line says "the desktop app generates a self-signed certificate that's unique to your machine and configures your system to trust it", so it's not the same as the Lenovo screw-up, but I had no idea it was there until today.

It probably ought to be, at the very least, an opt-in, don't you think?

Blizzard's Post:

Our recent update to the Blizzard Battle.net desktop app made sure players could properly use features like logging in to Battle.net via a social network, or joining a Blizzard group via an invite link. To facilitate these features, we updated the local webserver to use a self-signed certificate to be consistent with current industry security standards.

For those interested in more detail, using these features requires your web browser to communicate with the Blizzard Battle.net desktop app. Previously, the desktop app used a certificate signed by a public Certificate Authority, meaning that no modifications to your system certificates were necessary; however, this technique is incompatible with Certificate Authority policies and we can no longer use it.

While some browsers such as Chrome and Firefox are equipped to handle browser-to-app communication techniques, the changes were necessary for other browsers. For the time being, the desktop app generates a self-signed certificate that's unique to your machine and configures your system to trust it.

An important lesson and teachable moment thanks to more troubles with Google+

https://www.blog.google/technology/safety-security/expediting-changes-google-plus/

David Thacker / VP, Product Management, G Suite / Posted yesterday, under the title: "Expediting changes to Google+":

In October, we announced that we'd be sunsetting the consumer version of Google+ and its APIs because of the significant challenges involved in maintaining a successful product that meets consumers' expectations, as well as the platform's low usage.

We've recently determined that some users were impacted by a software update introduced in November that contained a bug affecting a Google+ API. We discovered this bug as part of our standard and ongoing testing procedures and fixed it within a week of it being introduced. No third party compromised our systems, and we have no evidence that the app developers that inadvertently had this access for six days were aware of it or misused it in any way.

With the discovery of this new bug, we have decided to expedite the shut-down of all Google+ APIs; this will occur within the next 90 days. In addition, we have also decided to accelerate the sunsetting of consumer Google+ from August 2019 to April 2019. While we recognize there are implications for developers, we want to ensure the protection of our users.

Details about the bug and our investigation

Our testing revealed that a Google+ API was not operating as intended. We fixed the bug promptly and began an investigation into the issue.

Our investigation into the impact of the bug is ongoing, but here is what we have learned so far:

- We have confirmed that the bug impacted approximately 52.5 million users in connection with a Google+ API.

- With respect to this API, apps that requested permission to view profile information that a user had added to their Google+ profile—like their name, email address, occupation, age (full list here)—were granted permission to view profile information about that user even when set to not-public.

- In addition, apps with access to a user's Google+ profile data also had access to the profile data that had been shared with the consenting user by another Google+ user but that was not shared publicly.

- The bug did not give developers access to information such as financial data, national identification numbers, passwords, or similar data typically used for fraud or identity theft.

No third party compromised our systems, and we have no evidence that the developers who inadvertently had this access for six days were aware of it or misused it in any way.

We have begun the process of notifying consumer users and enterprise customers that were impacted by this bug. Our investigation is ongoing as to any potential impact to other Google+ APIs.

Next steps for Consumer Google+

We will sunset all Google+ APIs in the next 90 days. Developers can expect to hear more from us on this topic in the coming days, and can stay informed by continuing to check the Google+ developer page.

We have also decided to accelerate sunsetting consumer Google+, bringing it forward from August 2019 to April 2019. We want to give users ample opportunity to transition off of consumer Google+, and over the coming months, we will continue to provide users with additional information, including ways they can safely and securely download and migrate their data.

A note for our enterprise customers

We are in the process of notifying any enterprise customers that were impacted by this bug. A list of impacted users in those domains is being sent to system administrators, and we will reach out again if any additional impacted users or issues are discovered.

G Suite administrators are always in control of their users' apps. This ensures that G Suite users can give access only to apps that have been vetted and are trusted by their organization. In addition, we want to reiterate that we will continue to invest in Google+ for enterprise. More details were announced in October.

We understand that our ability to build reliable products that protect your data drives user trust. We have always taken this seriously, and we continue to invest in our privacy programs to refine internal privacy review processes, create powerful data controls, and engage with users, researchers, and policymakers to get their feedback and improve our programs. We will never stop our work to build privacy protections that work for everyone.

The importance of THIS phrase: "We discovered this bug as part of our standard and ongoing testing procedures and fixed it within a week of it being introduced."

I want to formally congratulate Google for this... and I think it brings up an extremely and increasingly important issue: If a bug exists, or is introduced at some point, it can subsequently be discovered by any one or more of four possible entities:

A malicious blackhat hacker who leverages it to damage the company who created the software and/or their users or customers who have trusted the company's offerings to be safe to use.

As we've recently seen several times, it might be discovered by a self-aggrandizing greyhat who does not themselves leverage the discovered vulnerability, but who discloses the bug's existence irresponsibly to both embarrass the company and endanger their users or customers.

Or the bug might be discovered by a security-researching whitehat who responsibly discloses their discovery and quietly notified the company to that the problem can be repaired.

Or... the bug might be discovered by the company's own employees who can, as Google just has, immediately repair the oversight, determine the nature and extent of their users or customer's exposure, and notify them of any consequences.

While inadvertently "outsourcing" the discovery of unknown bugs to third-party security researchers is what often happens, it also often happens that those discoveries are made by those whose hats are not white. So it's a gamble. And, unfortunately, it's a gamble that many companies make because they do not, cannot, or choose not, to build-in processes to find those bugs themselves.

We know that those who create a system that's intended to be secure and bug free are the LEAST able to attack it and find fault with their own work. So some sort of "Red Team" facility needs to be created. It could be a separate group within the company who pride themselves on punching holes in the company's offerings, or an automated system which continually scans, fuzzes, and searches for defects and odd behavior.

The point is, the establishment and maintenance of any such proactive vulnerability discovery system is not free... and the value it returns to the company can be difficult to justify. One of the lessons we have ample opportunity to learn and relearn on this podcast is that security is not free. And companies do not make headlines for not having serious security incidents. They only make headlines when their security fails.

So I salute Google for clearly having processes in place -- which are working -- to often find problems themselves rather than relying, even inadvertently, for them to be discovered by others. We learned from Paul and MaryJo that Microsoft had much more than ample feedback from users of their most recent Windows 10 feature update disaster, which was warning of every problem that later surfaced at great cost to them and many of their users. So we know it's not enough to have a system in place... it must also be heeded when it's telling you something.

Way back 30 years ago, when SpinRite was young and I had an entire bullpen of technical support people on the phone with customers we learned that there's no such thing as a "one off." Anything that happens to one user will happen to others. But that's a lesson that takes some time to sink in. And since it's not good news, some people and some organizations prefer not to believe that particular lesson.

So... props to Google. In this regard they are a company worthy of emulation.

~30~