FTC PrivacyCon 2017
January 12, 2017
Segment 5
Transcript

MARK EICHORN: I want to thank everyone for the presentations through the course of the day. It's been a really interesting day for me. I'm Mark Eichorn. I'm an Assistant Director in the Privacy Division here at the Federal Trade Commission. And today we're going to be talking about infuser get infected by opening a malicious document, malicious attachment, or visiting an exploit website or a compromised website that is under the control of the attacker. Or, simply by opening an malicious binary.

This is an innocent-looking word file that a user may receive in an e-mail attachment. Once he or she opens t

sample. For example, in this case, we want to see how the cryptosystem works. Is it a new kind of ransomware attack, or is it similar to the other ones? If its new, what features does it have-- in order to generate signature, in order to do more research from a defense side.

So we proposed UNVEIL which is a sandbox. It's a behavioral-based approach. It sits in the kernel, which is in privileged mode. And it uses Windows kernel development technology like other antivirus products. Or, for example, backup solutions. They sit in the kernel. And we generate a fake but interesting environment for malware to run and sit in the kernel and look at their false system activity of the user mode processes. And we run the sample and look at how the sample interacts with the files and collect the corresponding information.

I'm not going to talk about what we collect, but there is a high-level table here. So this is a typical scenario. And on cryptowall attack, one of the current ransomware attacks, we collect I/O operation, which is the lowest level of-- is IRP, which is the lowest level of I/O operation between the file system and Device Manager. And as you see, for example, the process of opening files, reading their content, and encrypted with high entropy payload, which is an encrypted version of that payload. And then close the file. And as you see, exactly the same I/O operation is performed on another file, which shows that there is some kind of repetition that happens during an attack.

And as you see, svchost is a benign process, but attacker or malware authors, in order to evade detection, kill the original one, and just invoke a customized version in order to do those types of attacks. In order to evaluate our approach, we use the unknown samples. In the wild, we collected 1,200 samples every day. We had an infrastructure worth 56 UNVEIL-enabled virtual machines. We ran the samples in a parallel fashion and collected the results in order to do the tests, the analysis. We cross-checked our results with virus total, which is a collection of AV Scanner-- you can easily [INAUDIBLE] release, submit a sample and get a report about those samples. So in our approach, a new detection is that if a sample is labeled or detected as ransomware in our system, we submit it to the virus total. If none of the AV Scanners in virus total detected that as ransomware or any kind of malware, we label it as a new detection.

So we tested our system for a year with 148,000 samples, and this is the higher-level results of cross-checking with virus total. And we submitted each sample six times in order to see whether the detection ratio changes during this period. And we define a measure which is called pollution ratio that is between 0 and 1. The values close to zero mean that none of the AV Scanners were able to detect the sample, and the value was close to one means that all the AV scanners were able to detect that sample. So as you see, in the first submission, around 70% of the samples had pollution ratio equal to 0. That means that in 70% of the samples that we detected as

these are trying to deceive you and trying to get you to unintentionally install some of this unrelated software. And some of our work of this year-long investigation is to try and understand the relationship of unwanted software to these commercial pay-per-install programs, understand some of the deceptive promotional materials, and some of the negative impact that this whole ecosystem has had on users.

And so first, I'll kind of take you behind the scenes of commercial pay-per-install programs. And commercial pay-per-

programs. And so, Google tracked the numbers of how man

verify the site certificate that it is receiving from outside, also securely establish the TLS connection. And then it also has to send or issue a new certificate for your browser so that the browser can also be happy about it. And to make the browser happy and accept that new certificate which is actually issued by the TLS proxy in your AV, the Avi, the TLS proxy, must also insert its own root certificate into the system store which the browser relies on for trustworthiness.

Now, in terms of client and TLS interception, there are some advertisement-related products which also used to do it. And you can, of course, remove them. For anti-virus products, they generally do it to check web malware. Parental control applications-- they do it to check for

[APPLAUSE]

MARK EICHORN: All right. So we welcome questions from the audience. So come up to the microphones if you have a question. I will sort of summarize a couple of the themes that I took away. Obviously the papers are very different, so it's difficult to come up with themes that go across.

But the main one that occurred to me was that it's a challenging environment for consumers to secure your systems. Right You've got people looking to put malware on your system. You've got a whole industry looking to put nuisance or other unwanted adware or ad-injectors or scareware on your system. And then you've got this irony of, people, when they do want a tool to improve their security, there's this effect of sort of increasing their risk to man-in-the-middle attacks to some extent.

So the other thing that struck me is the theme-- not in all the papers, but in two of them. The first two were just, again, this theme of the technology arms race that we've been hearing about over the course of the day. Like in the previous panel about how there might be a safe browsing countermeasure to some of this software and then there's a countermeasure to the countermeasure. Or similarly with the ransomware-- the paper described certain techniques like the stalling technique and so forth to respond to maybe attempts to analyze the ransomware.

So why don't I dive in with some questions. And I guess I'll start with Amin. The one that jumps out to me from your paper is is this something that-- how will we implement this? Is this something that consumers could use on their side? I know you said that researchers can use this now. Is this something that a consumer would be able to use to sort of help keep their computers clean of ransomware, or is it something anti-virus companies would use?

AMIN KHARRAZ: So thanks for your question. So right now, it's specifically designed for malware research and reverse engineers. In practice, in order to defend against ransomware, there are at least three ways. One is educating the users. Users should not click on every advertisement they see on their browser. Users should not open any attachment they see in their inbox. Or they should use a good backup policy in order to make sure that the data will not get lost. Incremental backup is something that everyone is recommended in order to defend against ransomware.

So this is the most reliable approach to defend against ransomware for consumers. But there are also other ways to minimize the risk and also help the community, for example, that we are working on. One system that we propose actually defines concrete model to detect ransomware. This doesn't mean that current AV scanners cannot do that or don't have it. They may have it, but they don't-- they have it for other malware families, but the thing that they probably don't have is a concrete model to detect ransomware. Because ransomware is doing something in a very specific form.

And referring to your question, the approach that we have, by minimal modification to the tool, we can use it in the end user. This is actually the research that we are doing right now. It's a behavioral-based tool that sits in the user's computer and tries to find behavior that is similar to

ransomware. So there are some challenges, for example, usability is an issue. But we didn't have a huge overhead on this when the system actually is working.

But again you give some freedom now, you're at the behest of the curators of these stores as to what you can install and what you can't install.

MARK EICHORN: Mohammad, so if I'm a consumer going to choose between different anti-virus products-- some of them use proxies and some don't, right?

MOHAMMAD MANNAN: Right.

MARK EICHORN: So are people sort of aware? Is there any way for me as a consumer to know that this product uses a proxy and this one doesn't, just to compare between products on that basis?

MOHAMMAD MANNAN: So whether they do TLS proxying or not, generally this is advertised. So because this is a feature that you may be interested in. So AVs or parental control applications-- they don't hide it. It's kind of a selling feature for some. So you can check it. And in terms of how they implement it-- whether it is safe to use, you can see some of our reports, and then maybe choose which one is best suited for you.

But in general, I don't use any of these products actually. We test them, but we don't use them. So if you have to choose something, my general recommendation would be choose something that does not do it, because it's actually pretty complex to do it properly. Browsers that continually update it. Some of these companies that manufacture browsers, they spend a lot of effort to keep browsers safe. And a significant part of their effort goes into securing TLS. And if you rely on another product which is just a side feature for them, to be as cautious as the browser manufacturers, I really doubt it. So it's probably better to avoid products that are doing TLS proxying.

MARK EICHORN: Damon you talked a little bit about the sort of assistance that the PPI distributors provide to sort of allow the anti-virus detection. And maybe the chrome safe browsing avoidance, and so forth, to sort of facilitate this. But could you talk about the bad intent or not of the other players in the industry? Of the publishers and the advertisers? Especially, at some point, particularly the people whose software is being used as bait, have something that people want, right? Not really?

DAMON MCCOY: Again, this varies case by case. But a lot of times, it's very deceptive. At They'll try and package it as something like a flash update or something like that, where it's not even really a flash update of any sort. It's just completely deceptively marketed. Or they'll have the flashing pop-up that says you're infected with something. And they'll get you to install this.

Or they'll throw out a bait video or some kind. So there'll be some salacious video, and you go to try and play it, and they'll say, oh, you need to download this codec. So the publishers, the distributors, and the affiliates oftentimes have very sharply honed and very deceptive things. And so, oftentimes, they don't have a real, viable software product that anyone would want to download. It's just purely based on deception on their part. So we did some analysis of this to show that a lot of the affiliates are operating in a very deceptive manner.

So this is actually the main challenge that we have. It's a game. Companies come up with patches, companies come up with new services, but at the end of the day, we see more advanced attacks on exactly the same service or the same feature.

AUDIENCE: Well I don't necessarily mean signature-based detection, but the same way that you and your work have identified, say the I/O properties of ransomware-- can the operating system sort of incorporate identification of those same properties to try to head off ransomware, sort of in the same way? So basically the operating system now has the responsibility to do some of the detection. Or it can identify some of the same properties that you are able to identify in your work and basically help defend against.

AMIN KHARRAZ: At least, I'm not aware of that. And one issue is that once these types of systems get active, usability becomes something that these companies have to make sure of. So, for example, what happens in Chromium is that-- three or four years ago it was not like this, but right now you have automatic updates. It doesn't ask users to update their browser. It actually does it automatically, because the default assumption is that the user does not do it, for example, for many reasons.

And I think the main issue that operating systems may have with this, companies may have with this, is they have to do extensive experiments on the usability. And for the approach that we were proposing, it's like some other anti-virus companies or backup solutions that work right now on

This is Howard Beales.

about what they're going to do. The number that stuck out for me was if you approved each mobile app request for data, there'd would be 231 requests per hour.

That doesn't work. Where consumers who care about privacy are going to have to get that privacy is from products that help them protect it that are marketed as privacy protection products. And if people don't buy them, it's because they don't care enough. In the same way as if they don't buy any other kind of a product they don't care enough to be willing to pay the costs.

JESSICA RICH: Well hold that thought, because I want to come back to the arms race, which I think somewhat contradicts a little bit what you just said. That everyone's been talking about. Well, OK, we'll come back to it. But let me get the input on the consumer perceptions from Andrew.

ANDREW STIVERS: Sure, so yes, I would echo much of what Howard said in terms of yes, some people care, some people don't. I do think it's important, and I think, for me, as an economist, the sort of best way to determine how much people care is to look at their consumption decisions.

I think that this field gives us a particularly interesting and complicated problem for consumer choice for a lot of different reasons. One big one is that your choice about privacy is connected to all your other choices about privacy. So if you fool with the very intricate and detailed privacy controls on a particular very popular social website, that may absolutely have no effect on your overall privacy profile, because you share information in all sorts of ways. And it's going to be really hard for consumers to figure out what's the value to me of spending a lot of time investing in privacy in one dimension when there's 16 other dimensions that maybe I either don't know or don't understand or don't care about for whatever reason.

So it's a really complicated question to get a look at the market and understand what people's values are in addition to the heterogeneity that we see. And I always kind of go back to-- and I'm

And how important is it, under what conditions, and what time of life, what context? And I think the FTC has made some really important strides in that area to say, in Helen's words, of course privacy matters, but a contextual inquiry is really important to understand what privacy means in that particular moment, in that particular context and the environment.

But I think there's this really important question, which I think is part of what both Howard and Andrew were going straight to is if people care so much about it, why do they find it so very difficult to protect it or to act in a way that would seem to conform to their stated preferences? And I think that's a question that we, one, have lots of answers.

And I would point to the same article and say that not only has Alessandro in that article said it's complicated, but in that research and numerous other papers produced by him and others like Jens Grossklags who was here earlier today, have highlighted all of the ways in which information asymmetries, cognitive biases, the public good nature of privacy, architectures, default settings, policies undermine people's ability to get what they want from the marketplace. And I would point to another paper that was written by none other than Joe Farrell who held, I believe, Andrew's job.

ANDREW STIVERS: My boss's job.

DEIRDRE MULLIGAN: Your boss's job, right. And he was exploring the challenges to the market provisioning of privacy, and he was looking specifically at privacy conceptualized as a final good. Meaning that people are actually shopping for privacy. And as we all know, very few of us rarely do that. Usually privacy is an aspect of a good or service that we're shopping for. Or maybe an instrumental good that we're interested in.

But even thinking of privacy as a final good, he noted that there were some particular reasons to question whether the information and other conditions efficient contracting would hold in the marketplace for people to actually get good policies that were consistent with their expectations. He wrote that if consumers understood the implications of different up-front privacy policies, and the policies are truly effectively disclosed, including drawing consumers' attention, then the demand shift effect-- consumers could basically shop with their feet, because you could discipline market actors. There would be incentives for firms to choose responsible policies.

But given what we know about cognitive and informational barriers from Alessandro's work and Jens's work and others', that those prese

Consumer Reports activity might be ever so important is that privacy isn't a good that you can shop for on the front-

And it's inevitably going to be a back and forth, because there are smart pe sm saon(r)3(e)4(2(aor)3(e)4(2()-10a)

Were they actually choosing devices because they thought I don't want to pay an extra dime or $2 for security? No, they didn't understand that the things we're going to create security vulnerabilities that they were vulnerable to. But they might have if they were given the choice for the less secure or the more secure product. They might have chosen-- many people might be in the position where that extra $2 for the secure device is really more than they have to spend. But the fact of the matter is that those choices aren't going to be borne by them, because their devices were used to attack other people.

And so when we think about security for sure, and the negative externalities of the lack of security properties in devices that are proliferating and managed. Many of them don't even have interfaces for people to manage. That's a huge problem and it would be really, I think, damaging for all of us if we decide we're just going to let people make decisions in the marketplace, because we know even big companies don't always shop with security in mind, right? So I think we need to be thinking about, what are the minimal security properties that should be on those IOT devices? And surely updates is one of them, right? You should either be able to update it, or kill it, right? Because otherwise we're going to have lots of unmanaged, unupdated devices, that can be assembled into a nice little army to attack important assets. And that seems like a really bad outcome.

On the privacy side I think we're facing some really fundamental shifts, in that machine learning, right, is all about inferences. And if we assume that people are going to control their privacy by making decisions about what they disclosed to others, and it turns out that really benign data that you disclose can be mine to infer really sensitive things about you, the whole model of suggesting that individuals are going to be able to exercise control over who knows what about them, right, which is kind of essential to the fair information practices, doesn't really work. And so we certainly need better ways of giving consumers control over their information.

JESSICA L. RICH: Well a--

DIERDRE K. MULLIGAN: We're going to also need controls on what people can do with it. And you can look at something like the way we've handled genetic information, where we understand that the information doesn't just have implications for you. It has implications for other people that share traits with you.

JESSICA L. RICH: Well and adding to that complication is the idea that you don't even know all the companies that are collecting your information. Many of them are in the background. There's no more one on one relationship.

HOWARD BEALES: But that's the core problem with FIPS is that you can't think of this as a problem of controlling information. You can think of this as a problem of trying to control bad things that people might want to do with information. But if you think about controlling it as controlling information, that once you've given it to anybody can be passed, and in one of the papers I'll pass on all my friends information, because I don't care much about my friends. But the--

JESSICA L. RICH: Even me?

HOWARD BEALES: It's an information use problem. Not an information control problem. And

JESSICA L. RICH: Well, I'd like to add that some of the debate that-- and some of the research that we've been focusing on here-- does seem uniquely focused on a privacy regime that depends on consumers managing their own privacy a lot. Whereas, in Europe, they view privacy as a fundamental right, and there are more substantive requirements, not more enforcement, but more substantive requirements in the law. So a lot of the research today, and the debate around consumers managing their own privacy, is very uniquely American, I think, in terms of what we expect our consumers to do.

HOWARD BEALES: I would just say I think it's telling that in the United States privacy is largely driven by a Consumer Protection Agency, and that's really what it ought to be about, is consumers. It's not about data protection. Who cares? It's about protecting consumers from the

interesting findings in one of the early papers today was consumers care about the accuracy of the inference, about the accuracy of the prediction that's being made.

HOWARD BEALES: There's positive privacy rents. And negative privacy rents, right? If I'm the high value consumer, I really don't want greater accuracy. I don't want greater predictions, because I get a lower price, because you don't know I'm the high value guy. So I think there are some tradeoffs. But absolutely, there's more incentives to provide greater accuracy.

JESSICA L. RICH: Dierdre, do you have something to add?

DIERDRE K. MULLIGAN: Yeah, so I want to say, I think, first, I don't think that consumer protection is the only lens through which we should view privacy, right? It's an important human and political right. And while the Federal Trade Commission's activities have been incredibly important, I think that there are other ways in which privacy needs to be protected. And at times the activities that happen in the commercial sector have really important ramifications for government collection of information, as we found out very loudly and clearly during the Snowden revelations. On the, is it just the same? So certainly the data protection issues, and information privacy issues are going to be similar, but on steroids. And part of that is just that many of the devices that are proliferating, they have no interfaces. They don't provide any notice. There's no indication to consumers that data collection is happening. So it's very, very hard to use your exit, or your voice, right, to express a preference when you have no indication of what's happening.

But I think that there are other challenges that are going to arise, that I think are fully in line with the sorts of issues that the Federal Trade Commission looks at, which are things about, as we have all this data, and we start to be able to use it in real time, not just to alter prices, but to alter your experience of an environment to nudge you, and you know, concepts like undue influence, and manipulation. And you know, there's a huge potential increase in information asymmetries that I think are going to create an increasing set of issues for a functioning marketplace, that consumers are simply not going to know how it is that their experiences are being shaped. And so I think that there's another set of issues, that I think one can think of under the concept of privacy. We like to have this idea that there's some autonomous decision making happening, or some decision making that is free of undue influence, right? Or concepts people have talked about, fiduciary duties, and things like that, that I think are related to the increase in data collection, and the way in which it can be brought to bear on people's experiences, and life opportunities, that are going to be much more important than they have been.

MARK WEINSTEIN: And my question is really around this whole idea that privacy--

consequences, not the data, the problem is it's going to the government. It's not that the information exists, or that it was collected somewhere. And I just disagree completely with the premise about your grandparents, because I don't know where your grandparents grew up, but mine grew up in a small town, and everybody knew everything. There was no privacy.

MARK WEINSTEIN: In the privacy of their home, there was privacy. There wasn't a camera.

ANDREW STIVERS: So let me just put in a plug for the kind of authority the FTC has, which I think is really valuable, and an incredibly, incredibly powerful. So one of the advantages of doing the way we do things is that we allow the innovation to flourish in the marketplace. Generally, the kinds of things that we're concerned about, consumer protection, aren't health and safety issues-- sometimes they are-- but they're things that we can often kind of re-address after the fact. So we can allow the marketplace to evolve, and then when we identify practices that we think are harmful to consumers, we can go after that, and send signals to the marketplace that say, hey this isn't right. Taking the kind of [INAUDIBLE] regulatory approach, that I'm familiar with from my work. It has a whole raft of other issues associated with it, and especially when technology is moving very quickly, it's sometimes difficult to craft a regulation that strikes the appropriate balance between allowing new uses of data, and making sure that there isn't any harm there. For food safety, which again was what I was involved with before, that technology is a little bit more stable. It's still evolving. But it's a little bit easier, I think, to regulate that. Plus, you know health and safety violations are a little bit harder to wind back, if you allow them to occur. So there's some differences in what we're looking at here.

HOWARD BEALES: And I guess the other--

JESSICA L. RICH: One comment form each and then-- I've gone 10 minutes over so--

HOWARD BEALES: I guess, the other difference I'd point to from food safety is the germs may evolve, but the threats on the internet, and the threats to hack into internet of things devices evolve a whole lot faster. The regulatory process is just not fast. I think the only way we could build a camera that we could guarantee would never get hacked is it didn't take pictures. No. Other than that, we can't rule out all possibilities because smart people are going to be trying to figure out ways around that if there's something that they can do with it. And that's inevitable.

JESSICA L. RICH: Dierdre, comment.

DIERDRE K. MULLIGAN: So the Federal Trade Commission has done yeoman's work on getting industries attention to the security of their products. You know, I believe that they will continue to do that as they've already started in the IOT space. And there is an education process that happens as companies realize that they have to be paying attention to known vulnerabilities, and making sure that they don't build them in, et cetera. That said, I would agree, if anybody is going to be given authority around security in the commercial marketplace it should definitely be FTC, and to the extent that DHS would get authority, it would be around critical infrastructure. But who knows what the next administration will do? But you know, I would certainly fully support giving the FTC greater clarity around their authority in the security space. But I think what they've done has been both fully supported by the law, and has been very effective and important.

JESSICA L. RICH: But we're working with all those other agencies, and we're trying to strengthen the whole approach to data security across the government. So I've kept everyone really too long. One plug I want to make is that although we have taken great delight in

public. And so we have a dedicated mailbox, research@ftc.gov, and we hope that you share with us and we monitor it regularly. Thank you so much, to our panelists, and to all the panelists today, and for all of you for coming. Happy new year.