Listener Feedback #159

Description: Steve and Leo discuss the week's major security events and discuss questions and comments from listeners of previous episodes. They tie up loose ends, explore a wide range of topics that are too small to fill their own episode, clarify any confusion from previous installments, and present real world 'application notes' for any of the security technologies and issues we have previously discussed.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-387.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-387-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here, our Explainer in Chief. He'll talk about the Oracle Java fiasco and a whole lot more, plus answer 10 of your questions, with a bonus 11th question. It's all up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 387, recorded January 16th, 2013: Your questions, Steve's answers, #159.

It's time for Security Now!, the show that protects you and your loved ones and your privacy online with this man right here. He's the Explainer in Chief, Mr. Steve Gibson. I call him that because he's so good at taking technical difficult subjects and making them crystal clear. And he's done that on many, many shows. Today, though, this is a Q&A. So we'll get lots of little bits of crystal clear. Hello, Steve.

Steve Gibson: And a nice update on the happenings of the week. We've had an event-filled, over in the security industry, week that we'll catch everybody up with.

Leo: No kidding.

Steve: And there was an interesting tweet that I got from a listener and SpinRite owner whose Mac was saved, his SSD on his Mac was saved. And so I just sort of retweeted that, and a whole bunch of other people said, wait a minute, I thought you couldn't use SpinRite on an SSD and blah blah blah. So I thought I'd take a little bit of time to explain the problem. And in fact I sent you a picture that you can share at the appropriate time, showing a very distressed-looking solid-state drive. It's no longer the case that solid-state drives are as perfect as we would think, and I'll explain why that is also the case.

Leo: Excellent, Mr. G. And then we have 10 questions from you, the viewers.

Steve: We do.

Leo: Why don't we kick things off? We've got a GoToAssist ad, and I'll get to that before we go to our Q&A.

Steve: Okay. So the big news - and, I mean, it's big enough that you were even covering it on your weekend syndicated radio show, The Tech Guy. And I thought what was humorous to me was that the media picked it up as the U.S. government is recommending that every one uninstall Java. And, okay, well, CERT is the standard computer response team for the government that maintains a list of vulnerabilities and concerns. It is now a division of the Department of Homeland Security, the DHS. And for some reason what got caught in the news net was that the Department of Homeland Security is recommending that everyone uninstall Java.

Now, this is nothing different than any other zero-day flaws. And in fact I saw some people tweeting out to their followers and mentioning SGgrc, which is why I happened to see it, in their Twitter stream, that it's like, yeah, nothing new. Gibson's been telling us about this for years. So in fact this is nothing new. This is yet another zero-day flaw. It was discovered, I think it was just after our podcast last week, so like Thursday. And then the news continued to escalate. I thought, since it's easy just to say, oh, ho hum, I wanted to share from the DHS CERT notice what they said, just so that our listeners get a sense for how complex this stuff is.

So what the DHS posted was actually derived from the discoverer's and reverse-engineer's analysis, and that's our friend we've talked about before, Adam Gowdiak. He posted over on the secure list his analysis. This is, believe it or not, a simplification of what Adam posted. So CERT said - their title was "Java 7 fails to restrict access to privileged code." And we'll talk about version numbers and all that in a second. They posted:

"By leveraging a vulnerability in the Java Management Extensions (JMX) MBean components, unprivileged Java code can access restricted classes. By using that vulnerability in conjunction with a second vulnerability involving recursive use of the Reflection API via the invokeWithArguments method of the MethodHandle class, an untrusted Java applet can escalate its privileges by calling the setSecurityManager() function to allow full privileges, without requiring code signing."

Okay, now, that, believe it or not, is the simplified version of what Adam posted, which just, I mean, your eyes cross, and you can't uncross them. But, I mean, this stuff is really convoluted, but that's the problem is that we're now seeing an aging platform, Java, which has a massive code base. It was created in an era where security wasn't a focus. And it's been exploding in size as they've added major new features. And in fact this vulnerability didn't exist in 6 and earlier. There have been some conflicting reports, but the most authoritative sources say no, this is a consequence of new functionality which Sun/Oracle added at Java 7. So it's been all Java 7s up through Update 10, which was current until on Sunday, so just four days ago, Oracle put out an emergency patch in order to fix that.

So CERT goes on, saying, "Oracle Java 7 Update 10 and earlier Java 7 versions are

affected. OpenJDK 7, and subsequently IcedTea," which I guess that's a weak version of Java, "are also affected. The invokeWithArguments method was introduced with Java 7, so therefore Java 6 and earlier is not affected. This vulnerability is being attacked in the wild and is reported to be incorporated into exploit kits. Exploit code for this vulnerability is also publicly available. We have confirmed" - and this was the other thing, is this was - so this is all Java. This is all platforms.

So CERT is saying, "We have confirmed that Windows, OS X, and Linux platforms are affected. Other platforms that use Oracle Java 7 may also be affected. By convincing a user to visit a specially crafted HTML document, a remote attacker may be able to execute arbitrary code on the vulnerable systems. Note that applications that use the Internet Explorer web content rendering components, such as Microsoft Office or Windows desktop search, may also be used as an attack vector for this vulnerability."

Leo: And unfortunately it's pretty darn hard to turn it off in IE.

Steve: Well, okay. So the good news is that, with Update 10, which is not very old, because we were talking about this just a couple weeks ago, they added, they did add in the control panel an easy means for simply disabling the web aspect of Java.

Leo: When you say "they," this isn't Microsoft, this is Oracle.

Steve: No, "they" Oracle. And so...

Leo: And is it in the Windows Control Panel or in a special Java control panel?

Steve: Well, you get to it from the Windows Control Panel applet, which contains all of the control panel applets that have registered themselves. And so there will be a Java applet in the Windows Control Panel. It opens the Java control, and then there's a tab there, and a checkbox at the top that you can uncheck to disable Java's plugin capability across all of your browsers that are using Java. So they finally added that in 10. So even in this vulnerable version, that's there. Of course, the problem is only our listeners have gone in and turned that off, or uninstalled Java. And of course Oracle brags that it's in three billion devices.

Leo: Yeah, when you install it.

Steve: Billion with a "B," yes. When you install it and tell them you don't want McAfee to come along for the ride.

Leo: By the way, McAfee wouldn't help you in this case.

Steve: No. So our friend Adam Gowdiak, I'm skipping all of his technical stuff because it's, again, it's just - but it gives you a sense for you've got to be a deep expert in order both to develop an exploit and also to figure out what's happened and fix it. And but the

problem, the whole problem is this is a house of cards. It is incredible, ridiculous level of interaction such that even the authors of the code can't fix it. And what's really interesting is this is the result of a failed previous patch. Last time there was - and in fact our really astute listeners may remember some of that language because I shared it last time. This was when Adam found something like this in August of 2012.

So he writes this time, "This is not the first time Oracle fails to sync security of Core and the new Reflection APIs, just to mention the Reflection API filter. This is also not the first time Oracle's own investigation and analysis of security issues turns out to be not sufficiently comprehensive." And then he points to Issue 50, which was discovered in the code addressed by the company not long ago. Then he finishes, saying, "Bugs are like mushrooms. In many cases they can be found in a close proximity to those already spotted. It looks like Oracle either stopped the picking too early, or they are still deep in the woods."

So I'm glad for all of this because this level of attention increases the pressure and allows the budgeting of more focus on Oracle. This is always about money. This is always about economics. I'm sure the developers are wishing they had more time to, like, check new releases, to verify security, to run regression analysis. There's going to be some tension between the developers, who really want to do a good job, and the managers of them, who are feeling pressure to put out the new features that they announced at some prior conference.

So this sort of attention, very much like - we were talking last week, Leo, how Firesheep put pressure on the major social networking companies to bring up their SSL barriers much more quickly than they would have, if they ever would have, without it. Similarly, I mean, this is not good. But the fact that this escalated to you talking about it on your weekend radio show, and the news covering it as a U.S. government advising people to remove Java, that gets Oracle's management's attention and allows them to put their focus where it should. And...

Leo: They did an update yesterday. Did that fix this hole? Or was that...

Steve: Yes, yes. Well, okay, yes. Yeah, there was that, and then there was also a Microsoft emergency out-of-cycle update that we'll talk about in a second. Apple was very proactive on this. The instant this came to light, they used a new feature in OS X, which is their plugin blacklist, which is a file called Xprotect.plist, and they declared that the current version of Java and prior were unsafe. And so all Mac users were immediately protected, as soon as Apple became aware of this. This got pushed out to connected Macs. And of course the disconnected Macs were never in danger because this is all about networking.

Leo: So that's not an update. You just get that automatically.

Steve: Yes. Without doing anything, your Safari browser will no longer run Java, as soon as a problem like this becomes known. And Brian Krebs, our friend who's maintaining a great blog, KrebsOnSecurity.com, he did a great blog posting shortly after this, which was "What You Need to Know About the Java Exploit." But what I got a kick out of was that, when I went there to see how far down it was and to see if I wanted to get a link to our listeners, already there's another zero-day.

Leo: What a mess.

Steve: I mean, after this. So if you go to KrebsOnSecurity.com now, what you will find is "New Java Exploit Fetches $5,000 Per Buyer." And Brian writes: "Less than 24 hours after Oracle patched a dangerous security hole in its Java software that was being used to seize control of Windows PCs, miscreants in the Underweb were already selling an exploit for a different and apparently still unpatched zero-day vulnerability in Java." And then he links to - or, I'm sorry, then his blog posting has an excerpt from what the hacker posted. The hacker posted: "New Java zero-day, selling to two people, $5K per person."

Leo: Wow.

Steve: And the hacker wrote, "And you thought Java had epically failed when the last zero-day came out. I LOL'd. The best part is, even though Java has failed once again and let users get compromised, guess what? I think you know what I'm going to say. There is yet another vulnerability in the latest version of Java." So that is the 11 update, the patched one. "I will not go into any details except with seriously interested buyers.

"Code will be sold twice; it has been sold once already. It is not present in any known exploit pack," meaning that it's still valuable because it's unknown, "including the very private version of Blackhole," which is rented. You rent this particular Blackhole exploit pack for 5K per month. He says, "I will accept counter bids if you wish to outbid the competition. What you get? Unencrypted source files to the exploit so you can reencrypt as necessary. I would warn you to be cautious who you allow to encrypt. They might try to steal a copy. Encrypted, weaponized version, simply modify the URL in the PHP page that calls up the JAR," that's the Java file extension, JAR, "to your own executable URL, and you are set. You may pm me." So we're basically just always in Java zero-day vulnerability rolling forward mode at this point. It's gotten that bad.

Leo: My god.

Steve: Yeah.

Leo: Unbelievable.

Steve: So again, from Java v7 Update 10 and on, you can just disable Java in your browser, and of course that seems like a good thing to do. Or one of the tips that Brian has on his what you should know about the Java exploit, and I would recommend to people, go to KrebsOnSecurity.com, and you'll have to go back about six blog entries, blog postings at this point to find that one. But he's got a really nice page with tips about dealing with Java, one being to do the dual browser approach, which is completely remove the plugin and seriously disable it from the browser you use mostly, and then have a browser you don't use often where you go to trusted sites and have Java enabled, if you need it.

I mean, when I've talked about this in the past, I've received mailbag and tweets from people, for example, in Scandinavian, where for some reason all their banks apparently

are based on Java applets. And so they have to have Java enabled in their browser in order to operate. But the reason we see Java is that it's cross-platform. It's a serious programming language, a nice mature programming language, which codes - it compiles to a sort of a pseudo instruction set, a so-called "virtual machine," and then these virtual machines are hosted on different platforms, on different architectures, on Windows and Linux and Mac, and other places. Originally it was developed for set top boxes. And so that's a time-honored approach.

The problem is that, when they said, oh, wouldn't it be nice if you could use Java on web pages, this is not JavaScript, this is Java. So this huge, behemoth programming environment that was never meant to be exposed to the Internet got exposed to the Internet through the web browser, thus the problem. So it's really become a major feasting zone for, as Brian puts it, "miscreants."

Leo: I love "miscreants."

Steve: So turn it off, or uninstall it, if you know you don't need it. And if you do, and if you have a habit of poking around on the shadier side of the Internet, you really, for your own sake, need to come up with some sort of a strategy for doing that poking around on your iPad or on a tablet or on a safe browser that doesn't have Java and only use a browser with Java enabled when you're going to sites where you know you have to have it. And of course NoScript for Firefox allows Java control on a site-by-site basis, as well.

So, meanwhile, we talked last Tuesday, or we talked last Wednesday, about the prior day's second Tuesday of the month Microsoft Patch Tuesday event, where they offered, we'll remember, seven patches that fixed a whole bunch of, I think it was like 12 or 13 vulnerabilities. Those were MS13-001 through MS13-007. Well, Monday of this week they just released the out-of-cycle emergency patch that they didn't obviously have time to get done by last Tuesday. So this is MS13-008, and it does fix the remaining glaring problem that they had. That was with their MS IE versions 6, 7, and 8 zero-day problem, which their Fixit did not completely fix. This now does completely fix the problem.

So anyone who noticed, if they were shutting down their system, and they were told, oh, install patches and shut down, which is one of the things that Microsoft does now, or if your system rebooted, if that's the way you have it configured or whatever, it was because Microsoft was sneaking out this fix which I'm sure they would have done it last Tuesday if they could. And as it was they did respond overall very quickly. When we were originally talking about this, it was like a week before, and I was musing whether there would be time for them to get it into the next Tuesday cycle, and there wasn't. But they did get it out as quickly as they possibly could.

And on the topic of no matter how dire the warnings are, better critical infrastructure security does not necessarily follow. We have in the news, IDG covered the story and Network World among others picked it up, that malware was found infecting two U.S. power generating facilities through USB. And I shortened the story a little bit, paraphrased it for the podcast, but basically IDG News Service said that two U.S. power companies reported infections of malware during the past three months, with the bad software apparently brought in through tainted USB drives, according to the U.S. Department of Homeland Security. And this is interesting. This is an acronym I hadn't heard before: the Industrial Control Systems Cyber Emergency Response Team. So that's a new version. It would normally have US-CERT. This is ICS-CERT, the Industrial Control Systems. So I'm glad we have that. Maybe that'll help.

And they said, in the first case, the industrial control system at a power generation facility was infected with "common and sophisticated malware," apparently through an employee's USB drive. The name of the malware was not specified. The tainted USB drive came in contact with a "handful of machines" at the power generation facility, and investigators found sophisticated malware on two engineering workstations critical to the operation of the control environment. Investigators did not find malware on 11 other workstations examined. ICS-CERT recommended that the power facility adopt - I love this - "new USB use guidelines" - which we talked about three years ago - "including the cleaning of a USB device before each use." Yeah, that'll happen.

In the second incident, a power company contacted ICS-CERT in early October to report a virus infection in a turbine control system. About 10 computers were affected, ICS-CERT said. An outside technician used a USB drive to upload software updates during equipment upgrades, CERT said. The malware delayed the plant's reopening by three weeks. So here we're seeing reports, thanks to us having something like ICS-CERT. Notice that these various power facilities are remaining anonymous, probably to encourage them to report these problems rather than giving them a big black eye for doing so. So that's good. But, wow.

We know from having tracked down forensically the way we apparently, we and Israel, managed to get Stuxnet into the centrifuge control system in Iran. We know that those were off the 'Net and that the malware jumped the gap using USB drives. And so remember that the original viruses rode around on floppies using the so-called "sneaker net" approach. That's where they lived. And now they're living on USB drives which are obviously very effective because USB drives tend to be highly promiscuous in their behavior, no fault of themselves, of course. So anyway, be careful of your USB drives. And the good news is Microsoft awoke to this several years ago such that the operating systems are far more careful about just immediately running the autoexec file on a drive when you plug it in as they originally were set up to do. It's like, oh, won't this be convenient. Yeah, mm-hmm.

So there's a concern that might be big, but I don't know. And that is what's been reported as a zero-day remote preauth vulnerability in Linksys routers. A company, DefenseCode, whose domain is DefenseCode.com, reported several months ago that they had found a remote preauth vulnerability, meaning that they can get around login vulnerability, which gives them root access to Linksys routers. There's a video, and I created a bit.ly link for our listeners. It's bit.ly/linksys0day with a numeric zero, so linksys0day, all lowercase, in case people wanted to go look at it.

Now, what it shows is them connecting to the router over 192.168.1.1, which is the inside the network local interface. Now, while it's not good that there is a way for malware to attack your router from the inside, well, UPnP allows malware to do that with everyone's permission. So the problem would be if you had disabled Universal Plug & Play because you're security aware and smart, and something bad already got into your network, it could attempt to see whether you've got a vulnerable Linksys router and bypass your cleverly changed administrator password in order to access the guts of your router. And this little video absolutely demonstrates that happening. I'm assuming that this is not exploitable on the public interface, that is, on the side facing the Internet. If it is, we have a disaster on our hands.

Cisco has ignored, basically, this report. They responded that they have already fixed the vulnerability in their latest firmware release. But DefenseCode checked the latest version of the firmware for the WRT54GL router, which of course is a high-end powerful router that lots of people use, the WRT54G series. And the latest firmware, which was 4.30.14, was still vulnerable. And DefenseCode stated that the nature of it leads them to believe it

is probably widespread across many routers that are Linux-based and probably goes way back. And this is 70 million routers worldwide. So what DefenseCode has said is, "According to our vulnerability disclosure policy, the vulnerability details will be disclosed in following two weeks at DefenseCode.com, on Bugtraq and Full Disclosure," on the standard disclosure mailing lists.

So Cisco has burned up their time and apparently done nothing about this. Again, there's no reason to believe this is a vulnerability on the public interface. There's a lot of reason to believe it isn't because, remember, the first thing that packets hit when they come in from the outside is the NAT, the stateful NAT router. It's an effective firewall. It's very rare that we see a problem there. It's not impossible. My own spoofability test, the DNS Server Spoofability Test, it routinely crashes people's routers. And I'm constantly getting reports from people about - actually when I was developing that we were crashing our beta testers' routers, and so I figured out what it was I was doing so that the regular test wouldn't do it. But then I put in a "Does this crash your router" test for people who wanted to find out because the concern is there may in fact be an exploitable vulnerability from the outside that I stumbled on, but because I'm not a bad guy, I've never pursued it in order to figure out how exactly to turn that into an exploit. I have no interest in that.

But so there's some possibility that something is there which is a problem. It may just be that it's a form of overload that's crashing the router and not something that could be exploitable, at least in my case. But so I don't want to alarm people needlessly. I would be surprised if DefenseCode found something that was remotely exploitable. But they may, and they may only be showing it on the private interface, the inside interface, because they don't want to overheat the world about figuring out what they've done. We'll know in two weeks, and I will certainly update us. Cisco has been as lame as anyone could be about these problems. And in fact one of our Q&A questions we'll be getting to focuses on that with some details.

So anyway, just to - I don't think this is going to be a problem. The only vulnerability is if something already got in your system, in your network, that it could then take over your router. But again, that's true for pretty much everything. So Universal Plug & Play does that, as I said, with permission. So we will keep an eye on this and let everyone know.

I just tweeted - this is now we're in miscellaneous stuff - a fabulous review of 5v USB chargers. I know this is totally random.

Leo: And geeky.

Steve: It is so wonderful.

Leo: What are the differences between 5v USB chargers?

Steve: They're huge. I made a bit.ly, bit.ly/usbcharge. So, and Leo, check it out. I mean, the good news is there's a winner. There's a clear, clearly superior USB charger. Apple's is highly regarded by this guy, who is clearly an electrical engineer. Actually I got onto him because Simon Zerafa, who is a frequent tweeter and amazing finder of things, found a page where this guy shows how part of the logic of the old 6502 microprocessor can be reverse-engineered by looking at the exposed chip.

Leo: I have most of these chargers. So I'm glad to see that - but, wow. So there's a chip in here.

Steve: So anyway, so I commend our listeners, bit.ly/usbcharge. And it's a wonderful page. He takes, I don't know, maybe 12, and puts them through all kinds of tests. He tests the ripple, he puts them under load, he shows how well regulated they are, he tears them open and looks at the physical construction, and everything is rated and ranked. And the HP tablet charger is the clear winner.

Leo: Which you can no longer get.

Steve: Actually they're on Amazon for $25.

Leo: Is it a 10-watt charger?

Steve: Yeah, it's a high-power, a little better than five volts and a little better than two amps. So it's a beautiful charger.

Leo: And that's what you want in every case, even though your device may require fewer watts. An iPad requires 10 watts to charge. Actually the new iPad charge is even a little more powerful. But even if you have one that only requires, I don't know, .8 milliwatts or something...

Steve: Well, and as you know, power equals voltage times current. So five volts times two amps is 10 watts.

Leo: Right. But it wouldn't hurt anything if you put a 10-watt adapter on it.

Steve: Oh, absolutely not. It's just...

Leo: It would charge faster in most cases.

Steve: It's cruising along and having a good time.

Leo: Safety standards is one of the things he talks about.

Steve: Yeah, oh, and there's some really interesting stuff about the clone chargers. There were a bunch of clones.

Leo: They're not safe.

Steve: That are, I mean, oh. And he shows the pictures, like the clones of the Apple charger, side by side. And you can clearly see that not as much fit-and-finish care was given to the clone as to the original one. And then he talks about the insides. And the clones just behave miserably. I mean, there's just probably a diode in there. And it's like, good luck.

Leo: Good to know. Don't use cheap chargers.

Steve: Do not want to use a cheap charger. The last thing you want is for AC to get its way up that cord to your precious phone or tablet.

Leo: He also rates for vampire use, that is, how much power the charger uses when it's not charging.

Steve: Yes.

Leo: And the Samsung oblong is best in that regard. Only 19 milliwatts when it's not being used. To compare it, the iPhone charger's 195 milliwatts when not being used. But the iPad charger's 62. It's the next best trip. Wow, this is really interesting. I love this...

Steve: Isn't that nice?

Leo: I love this fake counterfeit, which looks pretty good until you see that it is - the Apple one is designed by Apple in California, and the cheap one, the fake one, is designed by California. That's the giveaway right there, yeah. I have a bunch of these Samsung oblongs because I have a lot of Samsung phones.

Steve: Oh, really, cool. Well, now you know that you can leave them plugged in, and they're not burning too much power.

Leo: Yeah. Wow, that is an interesting article. Who'da thunk it?

Steve: I thought that was a great page. So speaking of who'da thunk it, if you put up for our video listeners or watchers, Leo, that JPG which was tweeted to me last week, that is an image of SpinRite's drive map running on a six-month-old Crucial 500GB, so half a terabyte, SSD. And anyone who ever believed that solid-state drives were like memory will be disabused of that belief looking at that map.

Leo: Those U's are all unrecovered sectors.

Steve: Yes. Now, what happened was the person who was doing this, and I unfortunately don't have his tweet handle in front of me, he just tweeted to me, in the middle of the day, I think it was the late morning, that SpinRite was working on his USB drive, I mean, I'm sorry, his SSD drive from his Mac. I think he'd removed it from his Mac, and he was running it on a PC. I note that it's getting a very nice speed. He's showing it, the chart there, the map, SpinRite shows it as halfway through 500GB in about 20 minutes. So it would be about a 40-minute run for the whole half a terabyte, which is about right.

So but he said he was running it on Level 1. And I said, "Oh, you don't want Level 1, you want Level 2." He's a listener. And so he knew that SpinRite could be used to recover SSDs. And his system was trashed. It would no longer boot. Nothing was working. So he did this. So he restarted it at Level 2. The difference is both Level 1 and 2 are read-only, and that's the key. You don't want to run Level 4, which is exercising the surface, because they are exercising the bits and nominally wearing them out, not very much, but there's really no point to it. It's unlike the operation of a physical media with a defect where you're exercising the bits in order to find the defect. Here what you want to do is run Level 2 because Level 1 is not permitted to fix anything. That's why he was getting all those "U" reports of problems. But SpinRite had been prevented from doing any repair.

So he started again on Level 2, and I got a tweet from him about an hour later saying that his system was running again. So SpinRite did fix his 500GB SSD by removing it from a Mac. And he ran it, I guess hooked it up to a PC in order to run it and then put it back and was a happy camper. However, what he tweeted, he said, "Six-month-old Crucial 500GB SSD not looking good. KNEW I should have stuck with OWC."

So what happened was I retweeted his tweet, and it caused a flurry of responses from people saying, what? I thought you'd been saying don't use SpinRite on SSDs because it'll wear them out, and it won't fix them. And we have had now many testimonials from people like this guy who have had SpinRite successfully repair their SSDs. The reason they're failing is the same reason hard drives are failing. And that is, the nature of the world, unfortunately, is economics. Everything is economics. And engineers are being forced to operate right at the limit of technology. They could make SSDs that were much more reliable, and arguably the first SSDs were. They were very expensive, and they were not nearly as dense. But they were very reliable, just like the earlier hard disks which, because they were cramming fewer bits on the surface, the bits were bigger, and they were easier to find.

And so what's happened with SSDs is, remember that the technology is like a little tank. Think of it as a tank of electrons, and there's some leakage in the tank, just sort of natural migration of the electrons out of the tank, like a very slowly leaky bucket. Unfortunately the tank is a little porous. And the idea is that you either empty the tank to have it be a zero, or you fill the tank full of electrons for it to be a one. And everybody's happy. Except that the problem is to increase the density to get more bits in the same, like in a per given unit area. They've made the tanks smaller. So the tanks have a smaller capacity, which means there's less difference between a full tank and an empty tank in absolute terms.

And then they went even - to make matters worse, they went from using these cells, tanks, which only had a - which were only either empty or full. Someone said, hey, you

know, we could store more virtual bits per cell if we did multilevel tanks, that is, if we filled the tank a quarter full, half full, three quarters full, or full. Then what that would mean is there were four different potential fullnesses. And we know four states gives you two bits. So by using so-called MLC, multi-level cell SSDs, they've been able to double the density kind of for free, but not at the same level of reliability.

Every single one of my SSDs is SLC, single-level cell. They're not affordable, even today. I mean, they are really expensive. And the reason is that in order to get high-density SLC SSDs you need to use a lot more chips. So they're just more expensive to produce. But, for example, I'm using only SSL - I don't have any MLC SSDs anywhere because I just - I'm old school, and I realize that manufacturers are always going to push the boundary. They're going to make these reliable enough. So SSDs work sort of the same way today that hard drives do. They're reliable enough. But because they push the boundaries of the technology to get the density up in order to be as competitive as they can and to keep their costs as low as they can, we have problems like this one.

And once again, SpinRite can - you can do a read-only scan of an SSD anytime you want, and it is just as good for it as a read-only scan of a physical hard drive with SpinRite. Why? It's because, just like a hard drive, where I've said the hard drive only knows it has a problem reading a sector when it tries, it's not omniscient. It doesn't know anything about the sectors floating around out there that it hasn't visited recently. So it's only when it tries to read it that it's able to gauge whether or not it can, and how much error correction is required. And that's the other part of this. SSDs use error correction just like hard drives. Because these bits are flaky, they've unfortunately made them flaky by forcing them to have too much capacity in too little physical size. They're now relying on error correction to make up for the fact that the bits are coming back bad.

So just like hard drives, if you run SpinRite on an SSD at Level 2, it will do a read-only scan, and it may be terrifying the SSD controller. It will be showing it that it's got problems. And it will be relocating sectors of SSD in exactly the same way hard drives relocate sectors on the physical surface. And if there's a problem that the SSD controller cannot fix and relocate, then if you're on Level 2, that allows SpinRite to go to work and do its data recovery work, which is clearly what did succeed for this guy who tweeted me last week. So that's the whole story.

Leo: That's cool.

Steve: It is the case that, unfortunately, the passion for cramming as many bits in as small a space as cheaply, as inexpensively as possible - well, actually cheap in both senses of the word - is creating a problem. It's no longer the case that they are super - they're certainly reliable physically. The reason I have them in my laptops is laptops tend to get bounced a lot more than stationary systems. And so there I think a laptop really is the right, well, and for example, remember iPods used to have hard drives in them. That was nuts. And of course SpinRite was able to fix those, too, very often.

Leo: So but you're not recommending against them in desktops, are you?

Steve: No, no, no. I would say they are probably - I don't even really know if I would say they're more reliable than hard drives. At this point...

Leo: They're not less reliable.

Steve: They're not less reliable. I would say the benefit is they are non-mechanical, and that's a good thing.

Leo: And they're super fast. I mean...

Steve: Yes, and there you go. Super fast.

Leo: There's the benefit. I mean, they're massively speedy.

Steve: Yeah. To give you an idea, though, of how belt-and-suspenders I am, the latest server that I built for us has six OWC single-level cell SSD drives running in - I'm sorry, four, four drives, the highest reliability I could find. They are also, the OWC drives, they are massively over-provisioned. That's the term used for the amount of unused space available for relocation. They call it "over-provisioning." And then I run those four drives in RAID 6, which is full 100 percent redundancy. Any two of them can fail to read, and I still am able to read perfectly. So I just don't ever want that to be a problem. And you invest once, and then you're not having to make trips to the datacenter all the time.

Leo: I just got the new iMac that has - Apple's got this new thing they call a "fusion drive," which is two drives that look like one. One's an SSD, and one's a spinning drive. And it's a total of 3TB. It's not RAID. It's not an Intel technology that is similar. It's their own technology, software built into the operating system, Lion and Mountain Lion.

And what it does is kind of interesting, is it moves the most accessed stuff to the SSD, and the less accessed stuff stays on the spinning drive. And it's purely for speed. But having just installed this computer, so I haven't - it hasn't tuned itself yet. But I'm getting 290-plus megabytes/second write speed and 430MB read speed on this thing without having it optimized. And I presume it will get faster in the stuff that I need to be faster. So this is very close to an SSD drive. Direct SSD drive is very, very fast. SSD is even faster, four or 500 mbps.

Steve: Well, and it really makes sense for reading. Of course, SSDs don't write quickly.

Leo: As fast, right.

Steve: They read quickly. And the reason is that the only way, the way you write to an SSD is you drain all of the tanks throughout a region, and then you selectively refill them. And you also need to do that at a higher voltage so there's actually a voltage booster in the SSD that brings the voltage up to a couple hundred volts. And that's needed. Remember we've talked before that the way that the SSD works is that you pierce the insulation, and you squirt the electrons through an insulating layer. To do that you need a lot of pressure, and voltage is pressure. So the SSDs bring the pressure up,

squirt the electrons through an insulator, and strand them out on a little conductive blip, essentially, in the chip, and that represents a one bit. And the problem is over time the integrity of that insulating barrier breaks down, specifically because we keep squirting electrons through it. It just - it hurts it. Sensing the charge is easy. We're able to do that electrostatically. That's when we can read effortlessly, very quickly, and with no overhead...

Leo: Well, and it's random access. There's no seek time, either.

Steve: Correct. Correct.

Leo: But I have to say, you make it sound like the writes are slower. They're still faster than spinning drives. I mean, they're fast, the writes. This is the Other World Computing solid-state in my MacBook Pro Retina: 208 mbps write, 500 mbps read. That means I'm reading a gigabyte every two seconds. It's really fast.

Steve: Wow.

Leo: Really, really fast, yeah. So I only use SSDs in everything. I'm pleased to see how well that the fusion drive performs, given that...

Steve: Even when it's new.

Leo: Even when it's new. It hasn't optimized at all. And it's got the 3TB capacity, which is nice. So I will do a little ad, and then we will get to your questions. How about that, ladies and germs? How about it? And we've got 10 questions from our audience, our "listener-driven potpourri," as Steve puts it, episode 159. And we do this pretty much every other week, so that's how we got to such a high number.

Starting with Bob Iiris, a listener in the USA, who offers a minor correction, Steve. He says in Episode 385 you refer to the - and this was an offhand comment you made, by the way - the Bernoulli effect/principle used to float the hard drive's head off the surface of the rigid disk, commonly called the "Venturi effect." But you were actually describing the so-called "boundary layer effect." They're related in that the Bernoulli effect is occurring within the boundary layer. P.S.: With your audience, you just can't get away with anything, can you, Steve. I think you got it right because I think you started saying Venturi effect, and then you said the Bernoulli effect.

Steve: Yes. And I thought I would share Bob's note and that it is about the boundary layer, which is true. And what's interesting is, remember when there were Bernoulli drives?

Leo: Yes.

Steve: We went through that phase?

Leo: They weren't fixed disks, they were rigid disks.

Steve: Well, you mean they weren't rigid disks, they were flexible disks.

Leo: Right. That's what was unique about them. They did the same kind of floating head thing, but on a flexible disk.

Steve: Exactly. And what's really interesting is we think of floppies as spinning, well, those of us who do think of floppies spinning...

Leo: What's a floppy, Steve? I've never heard of that.

Steve: We old-school people. You think of them as spinning very slowly. And in fact they do spin very slowly. I mean, you can see them spinning around. If you ever looked at the hub of a spinning floppy disk, it's not going fast. What IBM determined, though, was that that's as fast as they could go before the head would start coming off the disk.

Leo: Ah. Oh, that's interesting.

Steve: Yeah. So it turns out that you actually don't have to go fast for there to be enough air pulled along by the spin of the disk in order to get under the head and for it to no longer be in contact. The floppy technology is more like the magnetic tape technology, which has to be in contact with the head. Hard drive technology, which is what the Bernoulli drives used, is deliberately a floating head technology. And I remember in the pictures that the Bernoulli Company had, that made the Bernoulli drives, they showed their head pushing down and the disk actually flexing under the head. There was an air bearing, and the head was deforming the disk a little bit as it flew, but it was not in contact with the head, or with the disk, which by the way was in a removable cartridge. These were large, high-capacity, I think 40MB, which, ooh, boy, back then, high-capacity cartridge technology.

Leo: Ooh. 40MB.

Steve: Oh, you'd never fill those up, Leo.

Leo: We've come a long way, baby. Jared in Australia raises some useful concerns about the CrashPlan cloud backup: Going around trying to find the top backup solutions, I decided to re-watch the Cloud Solutions episode you did on Security Now!. I decided to put both CrashPlan and Carbonite to the test. Unfortunately, CrashPlan requires Java to use, just to let you know. This itself made me switch back to using Carbonite. Even with limitations of its own with file type, at least we know there's no additional security associated, or security issues, I should say. Since Java is so down with patches right now, the question I always think when an application

requires such software is, is my data still going to be there? You know, there are a lot of apps that still use Java. Minecraft does. A lot of our Citrix apps do. That's not the same as Java in the browser; right?

Steve: That precisely is why I chose the question, yes. I wanted to - I thought this was great. First of all, what I discovered when I was doing the cloud backup solutions podcast was many of these multiplatform solutions are Java-based for exactly the reason we were talking about at the top of the show, which is you write it once, and you can run it anywhere. And for an application to be running on Java, exactly as you say, Leo, is completely different from the browser to invoke a Java applet when you visit a malicious website. So definitely worth keeping these separate.

These cloud backup solutions, while they're Java-based, they're not Java browser-based, and that's where the problem is. And so, for example, if you were to use the security control panel to turn off Java in the browser, then Java on your machine, for example using these backup solutions, still works just fine, and you're okay. And the fact that Java has security issues doesn't mean that the users, that is, the developers using Java are at fault for using Java. That is, if you see what I mean...

Leo: Well, but, yeah, but that's neither here nor there. I mean, it's not their fault. But if it's a security issue, you still wouldn't want to use it. But it's not the same. I guess the security issue is if somebody gave you a program to run on your system, and you trusted them, they could make a dangerous problem.

Steve: Yeah, right.

Leo: But that's any program.

Steve: Exactly. An EXE is far more powerful than a Java virtual machine program.

Leo: True, good point. That's a very good point, yeah.

Steve: So I think it's...

Leo: So if they're going to give you a malicious program, they're not going to give it to you in Java. They're going to give you an EXE. They don't even need an exploit. They just say run this.

Steve: And none of our users would be able to use it if it were in Java because they've all uninstalled Java, if they had it in the first place.

Leo: But that is the disadvantage, also, to uninstalling Java is that you have that issue.

Steve: And also, if you think about how detailed the exploit was that I read, you could see that it was like winding its way through a circuitous route to get to a particular function where it was able to trick it into changing your security settings that then gave you remote access that you wouldn't have otherwise had. So here again, this is, as I've discussed before, Java allows all these things, and then they try to put a barrier up to prevent you from getting to them. That's fundamentally error prone. You don't want power which you then have to firewall. You'd really rather not have that capability at all. Then there's nothing, there's no possibility of figuring out a way to get around what's blocking it. Unfortunately, Java is a powerful environment. But these exploits are specific to remote browser-based implementation. And that's important to understand, too.

Leo: Very good. Question 3, or, yeah, 3. Mike King on the eastern shore of Maryland notes that helium-filled hard drives may lift Western Digital to the top. Have you heard about this, Steve? And either way, what do you think? It's a Computerworld article about hard drives, not with - I guess hard drives have, what, air inside them?

Steve: Yes, they normally run at atmospheric pressure. There's typically a little multi-hole filter in the lid, like in the cover of the hard drive. If you look somewhere...

Leo: To equalize pressure.

Steve: Exactly. And so air is allowed to pass both ways through that in order to keep the pressure equalized.

Leo: Otherwise they would, like, wouldn't work well at altitude and things like that.

Steve: Yeah, exactly, well, because, if you were at high altitude with low outside pressure, then the lid would bulge, and lord knows what effect that would have.

Leo: That wouldn't be good. If you think yelling at a drive messes it up, bulging lid's got to be a problem.

Steve: So very, very clever. Western Digital realized that the density of helium is - I think it's one-seventh the density of air. That's why helium balloons rise and go up in the air is that helium is so much lower density than air. They realized, if they switched to a helium environment, everything changes. It turns out that a substantial amount of the power that a drive uses is just overcoming the air friction of the disk platter surfaces as they're spinning.

Leo: Wow. No kidding.

Steve: And so the motor is working to spin the platters at a constant speed against the drag of the air. And if you switch to helium, it dramatically cuts the drag on the disks. It also completely changes the physics of the head/disk interface, as we might imagine. We've been talking about the boundary layer and all that. And it turns out that that

allows them to increase the density further because it allows them to fly the head closer, at higher speed, than they've ever been able to. So they could spin faster, get the head closer, and a closer head means a greater signal coupling between the disk and the head, and that means more bits, baby. So anyway, this is very cool. They've got patents on it. They beat Seagate, and is it - there's one other.

Leo: It's pretty much those two.

Steve: It's Western Digital, Seagate, and there's another foreign manufacturer.

Leo: Well, there's Hitachi. But, no, Hitachi is Western Digital.

Steve: No, Hitachi got bought by, or they bought IBM.

Leo: No, it says in fact that Hitachi was the ones who developed these helium-filled drives for Western Digital. That's why it's so confusing. Hitachi bought IBM's Diskstore business.

Steve: Right, the Travelstore line.

Leo: Yeah, Travelstore, Diskstore. And then Western Digital bought Hitachi. I think it's shaking out to really only a couple, and there might be some, like, small ones around.

Steve: Anyway, I just thought that's very cool, helium-filled drives.

Leo: It is great. It's great.

Steve: And when the drive died...

Leo: Western Digital has 45 percent of drives; Seagate has 48 percent. So there's somebody who has 7 percent. I don't know who it is. But they're not important.

Steve: Yeah, and when the drive dies, Leo, you could crack the lid and suck out the helium.

Leo: [High voice]

Steve: Exactly.

Leo: There's a guy who calls the radio show and does that every time.

Steve: Oh, goodness.

Leo: It's like, grow up. It's funny once. Luca in Verona, Italia suggests a possible UPnP solution: [Italian accent] Steve and Leo. I'm not going to do that. I just finished playing through Episode 385, and I thought of a possible solution for the listener asking about the PS3 and UPnP. He could get a second router and enable UPnP on that, daisy chain it to the first one, then place it in the first router's DMZ. This way he could have only the PS3 connected to the second router, which has UPnP capabilities, and keep the rest of his machines safe under the current network. He'd also be protected from any hypothetical attacks exploiting vulnerabilities in the PS3 since it would be also isolated from the rest of the network. Keep up the good work. Luca.

Steve: And that's very clever.

Leo: You've thought of that. Wait a minute, isn't it that triangulation, router triangle thing?

Steve: Yup. We've covered the idea of using routers as, like, one-way valves. And the idea of putting them in series, but connecting them...

Leo: But you need three to do what you wanted to do.

Steve: Well, you really do because, technically, if the PS3 were infected, it has visibility into the network outside of it, which would be your private network. So you really need a Y connection with one router for the PS3 and another router for the local network. That way both of the interior routers are feeding to a third public one. But he is correct, and I hadn't thought of this, that you could put the router that requires UPnP to be enabled, you could set that router's IP as the DMZ for the primary public router, so that unsolicited traffic would automatically be routed to that router, where UPnP would have been dynamically opening ports to allow access to the PS3. Meanwhile, your second internal router would have UPnP disabled and would essentially be barricading itself against both the PS3 and any mischief it might get up to and the outside network. So it's not lightweight, but routers are pretty cheap these days. So, yeah. Anyway, I thought that was clever. That's a variation I hadn't thought of before. So thank you, Luca.

Leo: Bill Burlingame in Huntsville, Alabama, has a question about the Quiet Canine project: Steve, I heard your update on the unit you're designing to stop dogs from sustained barking, No. 385 for those of you looking for that story. Is there a chance that, once the design is completed, a kit can be made available? Has this guy not been paying attention? If not, I hope the list of parts are readily available to people who don't happen to live in a large metropolitan area. My first choice would be a unit that is already built. If you don't have the time and resources to offer one, would

you grant permission to someone to sell units based on your design?

Steve: So that was a new twist that I thought I would share. First of all, you and I did talk about my intention to open source, open design everything. There is a company called Digikey which is fantastic, and all of the parts are sourced from there, and I've got a parts list with all of their own numbers. So it's sort of a virtual kit. And one of the people who is participating over in the Google group, it turns out that Spark...

Leo: SparkFun.

Steve: SparkFun has a division which produces PC boards at low volume. What they do is you're able to get a PC board that - normally the problem is you just can't get one little PC board because the PC board fab people want to fabricate them in huge sheets and then divide them up. Well, what SparkFun's company did, I think it's called BatchPCB, if I remember, I think that's the name of the company.

Leo: A clever idea.

Steve: BatchPCB. So the idea is everybody, there's like all these projects there. Anyone who wants a PC board puts in an order. It's not fast. You have to wait some number of weeks, typically. But even that's not bad turnaround for PC boards. And what happens is, as soon as they get enough of them for the total real estate, they submit that as a single composite board, which is then broken up into pieces. So it's very clever. And so the bottom line is that, one way or another, the answer is yes, Bill, there will be a way for somebody who can build these or knows somebody who can - hey, do me a favor, build this for me - to essentially have a kit. I really, myself, don't want to be in the hardware business. I also don't have any problem if somebody else wants to build them and sell them. And if they do a good job of that, I could certainly refer people to them.

I am going to still make a bunch of beta test units myself for our listeners who have this problem because I want to acquire more information about whether and to what degree this whole approach works. So, yes. Again, it's still premature. People write to me telling me that the Quiet Canine pages are all blank except for one or two, and it's like, yes, I know. That's because I don't have anything to put there yet. We're working frantically over in the Google group, so much so that people have asked how do I unsubscribe to this because there's just too many postings. It's like, I understand.

Leo: Interesting. Question 6, Mack Morris, Columbia, South Carolina. He's worried about the CBC information leaks. Steve and Leo, thanks for the show, blah blah blah. I have a question regarding cipher block chaining. If - oh, boy. Oh, boy.

Steve: Yeah. Just read this and don't try to understand it while you're reading it.

Leo: Good, because nothing will come through. But I am going to read it as if I understand it so that you can understand it.

Steve: That's good.

Leo: If I understand correctly, the first block of "N" bytes is XORed with an initialization vector. The resulting bytes are used to XOR the next block. This continues until all blocks have been XORed. If the key size and block sizes are different, wouldn't this result in information leakage? For example, if the key is 8 bytes and the block size is 10 bytes, then the remaining 2 bytes would be the same after XORing is complete. I'm assuming that the remaining bytes are XORed with zeroes. That may be a false assumption. You could just wrap around.

Wouldn't anything that is XORed with the zeroes result in the same bytes as the input? If the key is larger than the block size, the remaining bytes of the key would always be present at the beginning of the ciphered text. Well, that'd be bad. If the block size is larger than - I suspect this has been solved, but we'll see. If the block size is larger than the key, then the remaining bytes of the plaintext will be present at the beginning of the ciphertext, revealing the length of the key. While the length of the key may not be sufficient for the entire cipher to be compromised, it is information leakage nonetheless. Additionally, the first few bytes may contain information that does not need to be plaintext. If this is correct, would it be a good idea to pad the smaller item with ones instead of zeroes to prevent this leakage? Or is this how it's done anyway? Thanks, keep up the good work, Mack Morris, Columbia, South Carolina.

Steve: Okay. So there's a couple problems with what Mack has described. The reason we do anything called a "cryptographic mode," which is different than a cipher, is to keep information from being leaked. With a cipher, the actual algorithm, it takes some number of bits of plaintext and converts them, sort of maps them, into a completely different and unpredictable same number of bits in the output. But the problem is, every time you put the same thing in, you get the same thing out. So if all we did was encrypt a file by enciphering blocks of bytes, then patterns could be seen. That is, there would be information leakage because someone looking at the ciphertext might see, oh, look, this block is the same as that block, which means that the plaintext must also be the same. So while it doesn't disclose what the plaintext is, it tells you that it's the same.

And there are, for example, protocols that are very structured, where there's certain, like they have headers and trailers and so forth. They contain structuring information. For example, a JPEG of a certain size, it'll have in its header descriptions about the type of encoding, that it is a JPEG, the height and width and so forth. Those things could be figured out. You would then know what the plaintext was that corresponded to the ciphertext. And if you found that same ciphertext anywhere else, then you would know that it just happened to be the same plaintext again.

So what we need to do is we need to prevent that case, where the same input data, input to plaintext, results in the same ciphertext output. The way we do that is we, in addition to having a key - and remember the key is what determines the mapping between plaintext and ciphertext. Additionally, we have something called an "initialization vector," or IV. The initialization vector is always the size of the block, that is, the size of the cipher. Like AES is a 128-bit cipher, so the initialization vector is 128 bits. That's chosen at random. Doesn't matter what it is. And in fact it can even be in public. It can be known.

So that gets us started because we XOR the plaintext with the initialization vector, then we encrypt it, then we take the encrypted output of the first encryption and use that to

XOR the plaintext of the next block. And then we encrypt that to get the second block. We take that second block and use that to XOR the plaintext for the third block as we encrypt it. So that's why it's called "chaining." We're creating a chain from one block to the next, taking the output from the previous block and using that to scramble the bits going into the succeeding block. And so the beauty is that it creates an ever-changing scramble.

And remember that a good cipher for an input is it gives us a pseudorandom output. It is absolutely pseudorandom. It has no way of knowing, if you don't know the key, what input gave you what output. And so really smart cryptographers have thought about this a lot and said, yup, there's no way to gain any information that we don't want given. So that's how CBC, cipher block chaining, works, and why we need it.

Leo: I thought it was the Canadian Broadcasting Company. Tim in St. Louis has already got a hot flash drive: Steve, I'm sorry I'm a bit behind, but listening to 381 where you discuss a new flash technology that creates heat to prevent the drive from wearing out. On Black Friday I bought a new flash drive from PNY that was on sale. Unlike my existing drives, I noticed this one in particular gets extremely hot to the touch, almost too hot to touch, actually. I assume that particular new technology is not out in the wild yet from the way you spoke of it, yet this drive is hot for some reason. Even if it weren't designed for that purpose, could I have just gotten lucky in getting a hot running drive? Is there a chance it could prolong the life of the drive, or is heat like that not the same? Also, would keeping the drives in a warm area be beneficial? Thanks for keeping us informed. Tim.

Steve: Well, that's an interesting idea, Tim. I think you should probably send it back.

Leo: But I have a USB key that's quite good, it's very fast, and it's USB 3.0, and it gets fairly hot. I mean, I don't think it's unusual.

Steve: No kidding. Wow. I've never noticed heat. I've never noticed heat from mine. But for what it's worth, we're dealing with, okay, what I was referring to is a new heat-annealing process.

Leo: Right, that's different. That's not the drive getting hot.

Steve: Yes. And it's spot annealing at the molecular level with extremely high heat, like 600 degrees Fahrenheit, so way hotter than just power running through the flash drive and warming it up as a consequence. So I don't think, I mean, I don't think the level of heat you're going to get will be helping to anneal the flash drive. I would worry that it's going to hurt it over the long term. But they're robust against temperature, so that's probably not a problem. But no, the answer is we're a different level of heat. It's hundreds of degrees hotter, many hundreds. And also done sort of like on a microscopic instantaneous level, where nothing else has a chance to melt before the annealing process gets done. And we're probably several years away from seeing that in production. But, boy, will that be nice.

Leo: Sorry about that. I was doing sit-ups. Don't ask. Don't ask why. Steve, I only got two. I was trying to do 30. So take your time on this answer.

Steven Knight in Brisbane, Australia wonders about ransomware: Steve, over the Christmas break the Australian media covered ransomware, where your data is encrypted, and you have to pay a fee to get access to it. Or not. This is probably the mafia, after all. The media coverage didn't really get into prevention. My first thought was effective backups. You could wipe the affected device and restore the data, as well as firewalls, et cetera. Can ransomware be detected by traditional scanning software, and what are the best steps for prevention?

Steve: Okay. Leo, you do sit-ups. I'm going to answer the question. So ransomware, I think, is a really clever idea. I'm not endorsing it, of course. It's bad. But it's interesting. If a virus gets in your computer, then the question is, what is it able to do? I mean, normally it's able to use your machine as part of a botnet in order to DoS other sites. Maybe it watches what you're doing in the hopes of being able to see a credit card transaction or intercept banking work or something. But the other really interesting idea is that it could encrypt your system. And which is really an interesting sort of Catch-22, I mean, if it just wanted to hurt you, it could just reformat your drive. It could wipe out your files. But then it doesn't have anything. If it wants to extract value from you, what better thing to do than hide your own files from you, ransoming your access to them by making a payment. So when this surfaced a few years ago I thought, wow, that's diabolical.

Leo: Isn't it clever? Yeah.

Steve: Really clever.

Leo: I had a call over the weekend by a guy, he had ransomware. It says it's the FBI; right? They're smart. They're getting smart. They have a big FBI logo and stuff. But then it says that you have a fine, you have a $300 fine, and we will release it. Your data has been blocked because of a violation of the Patriot Act or something, and we'll unblock it. But then they ask for the payment in these cards, these prepaid cards that you go to the 7-Eleven to get. And I asked the guy, do you think the FBI uses prepaid cards for fine payment? Does that sound - MoneyPak, that's the name of them, MoneyPak cards. In fact, that's why they call this the MoneyPak Trojan. It's ridiculous.

Steve: So the bad news is, it is, in all other regards, it's malware like any other. You don't know where you got it. You went to a website. Something happened. You had Java enabled on your browser. There was a JavaScript exploit. You downloaded a bad piece of software, and it had a trojan in it, who knows what. One way or another, you've got this junk on your system. And as we discussed, antiviral software is having an increasingly difficult time because the code is encrypted. Sometimes it's polymorphic encrypted so that every instance of it looks different. It's very difficult for software to inspect the outside and see whether it's a problem.

Leo: This is good. I just pulled up the warning: "FBI Anti-Piracy Warning. All activity of this computer has been recorded. If you use a webcam, videos and pictures were saved for future identification." And then it gives you a barcode. And then it says your IP address. It says: "Your IP address and hostname were recorded for future identification. Your computer has been locked. Illegally downloaded material, MP3s, movies, or software have been located on your computer." Then it quotes U.S. copyright code, possession is punishable, et cetera, et cetera. "To unlock your computer and avoid other legal consequences, you are obligated to pay a fine of $400. Payment of the fine is done by Green Dot MoneyPak payment voucher. Failure to comply with FBI anti-piracy warnings could result in criminal charges and possible imprisonment of up to three years in country jail." So I think country jail is probably the giveaway that this - if the MoneyPak didn't get you, the country jail might get you.

Steve: I wonder if you could be in city jail, or if they just say...

Leo: Oh, no, you've got to go to country jail, my friend. And you know what that means. Oh, lord. But of course they need to anonymize the payment; right? Unless - actually the best thing to do would be get a credit card number, and then they could just keep charging you. Some of them do that.

Steve: Well, that's why the traditional hack was to use Western Union MoneyGram. That was the way the Russian mafia was always extracting money because it was a wire transfer. And baby, when that is gone, it is gone. It's over.

Leo: Yeah, no way to recoup that one.

Steve: Yeah. So unfortunately there isn't anything, Steven, that is special about ransomware. It's just a particular breed of very clever sort of late-model attack which leverages the malware's presence in your machine in an interesting way. It doesn't just wipe you out because it wants money from you.

Leo: But I should say that a backup is always kind of the last resort protection against corruption of any kind, including malware. If you do have a good backup, at least you can wipe and reinstall, which is really the best way to get rid of this.

Steve: Yeah. I would say, if it's easier for you to use Grandma's Green Eye Money Card, then go ahead and try that. [Leo laughing] And if that doesn't work, then maybe just go from the backup.

Leo: Mike Robinson in Michigan updates us one year after Reaver and Linksys WPS. Has it been a year?

Steve: Yes, Leo.

Leo: A year.

Steve: Yes.

Leo: This is of course a hack that allows you, even if you're using a secure WPA2 password, allows somebody to break into your system trivially.

Steve: Four to 10 hours, and typically half of that time, somebody just doing a passive network sniff is able to obtain access to your router.

Leo: Nice. Steve and Leo, I think Linksys needs another black eye. It's been over a year since Reaver - which is the tool that allows you to do this, free, open, widely available - was developed and released. And most of Linksys routers have yet to be updated. Remember Linksys routers, even though they had a Disable WPS button in some cases, it didn't do anything. It just pretended.

Steve: Now, and click that link and look at Cisco's list of routers which are still just flapping in the breeze.

[homekb.cisco.com/Cisco2/ukp.aspx?vw=1&docid=3bccc46248f9417b909e2c1028f6778e_WPS.xml]

Leo: To their credit, they posted this list quite quickly. Oh, too bad. They must have known we were about to check it. They say "We are currently performing system maintenance in the knowledge base. Please try again later."

Steve: No kidding. Oh, wow. For what it's worth, there's, like, maybe 50 routers, and I think maybe a third of them are not vulnerable. The rest are TBD, to be determined, when...

Leo: They're not going to update them.

Steve: No. No, you're right, they're not going to do anything about it. They're just going to - just like, oh, well, not a...

Leo: It's too late. If we wait long enough, they'll be obsolete.

Steve: It's a consumer product. Good luck.

Leo: Really goes to show, he goes on to say, that Cisco doesn't care all that much about the security of its customers. I have a WRT54G2 V1. That's a pretty old

router. I gave up waiting, so I flashed it - the nice thing is it's an old router that's easily flashed - with DD-WRT thanks to Know How, our show. We did a whole Know How episode on flashing your router. Frankly, you've got much better firmware on there anyway. Tomato or DD-WRT are much, much better and don't have WPS at all. There is still no - I didn't realize it had been a year. When you talked about Reaver and WPS, you pointed out it would be an almost trivial thing to fix. But what they're doing is they're just disabling WPS; right? Nobody has yet put out a WPS that's safe.

Steve: Correct.

Leo: Can you believe it? Just stupid.

Steve: I know.

Leo: Stupid, stupid, stupid.

Steve: And, I mean, what this tells us is there has to be massive pressure on a company before it'll overcome their inertia to fix it. It's got to be a big black eye. Which is why Oracle rushes out a Java patch, Microsoft...

Leo: Department of Homeland Security has to make an announcement. And you have to put it on the NBC Nightly News.

Steve: Yes. Meanwhile, Reaver is sitting there as a proof of concept. There's plenty of code now around. And bad guys can sit somewhere with secured WPA encrypted routers and, in a few hours, gain access.

Leo: Yeah. Question 10 wraps it up today. Peter Smith in London having trouble with GoDaddy and the "new" 5.0.0.0/8 public range. We talked about that a few episodes before. I've been a listener for over three years now. I enjoy the podcast greatly. It is enjoyable and informative, with enough technical information to make it interesting and informative without swamping a listener. He's obviously a pretty high-level listener. I get swamped almost every episode.

We have recently been allocated our own public IP address range from RIPE within the 5.0.0.0/8 range. This was allocated to RIPE in November 2010, and I believe they started giving those to customers in 2012. As you remember, this was previously used by Hamachi for their LogMeIn and other VPN solutions. It's kind of the other way around. LogMeIn used it for Hamachi and other VPN solutions.

We've run into issues since using the IP range to source NAT our office connections as it seems probably that people are blocking this with an IPS system. Both GoDaddy and Adobe are unreachable when I use an IP address in this range; however, when I switch the NAT to another IP address, I can connect. Oh, that's interesting. Obviously GoDaddy is used to host many sites and servers, and as a

result don't have access to any of them. Interestingly, with GoDaddy, ICPM traffic is unimpeded, but both UDP and TCP are stopped before the first hop on their network. I have contacted - ICPM is ping; right?

Steve: Yeah. Actually that was a typo from him. ICMP.

Leo: MP, right, yeah, yeah, ICMP.

Steve: Internet Control Message Protocol.

Leo: I have contacted both Adobe and GoDaddy, and they are currently investigating the issue. I'm currently interested in finding out which IPS systems they are both using - this would be a firewall or an intrusion protection system - since this is probably causing the issue. However, I think it likely they might be reluctant to give out this information. Obviously this is quite frustrating, and I would like to know what steps you would take if it were your IP range, especially as I think that it is likely this issue will pop up repeatedly for us. Peter. What an interesting point.

Steve: So this is actually not unexpected. The back story our listeners, our frequent listeners will know is that, as the Internet began to run out of IPv4 space, the traditional 32-bit addressing scheme, which uses the dotted quad approach, a previously unallocated block, which was huge, 16 million IPs all beginning with 5, were taken off of the shelf and put into service. This was the private - it was a private IP range only because no one had used it, very much like the 10-dot range. But 10-dot is officially for private networks. Five-dot was simply not yet in use. So what happened is overeager router configurers, that is, human people who set up Internet routers, and I've seen this many times in routing tables, they will block the private address ranges. They'll block 10-dot everything, 172.16 through 172.31, and they'll also block 192.168 because those are known to be private, meaning that no public router ought to ever see packets with that addressing. And somehow they get out on the Internet. I mean, the Internet's a crazy place. If it can happen, it will. And so the point is there's really no one to send that stuff to. So routers are often set up to just drop those.

I have also seen tables which drop all of the previously unallocated ranges. Some engineer, some Internet engineer somewhere thought they were being clever by saying, well, 5-dot is also unallocated. I'm going to block that because there should never be any 5-dot traffic. Well, that was once upon a time, but it's not any longer. So unfortunately it's probably not the endpoints, firewall, or intrusion protection system, the IPS that he refers to. It's lord knows what. It's a router somewhere out on the Internet.

Now, it is possible with various traceroute programs not to use an ICMP message, which he says has no trouble getting through, but you can use TCP or UDP protocol in a tracerouting fashion simply by changing the TTL, the time to live, making the TCP or UDP packet expire while it's en route. And that way you could determine which router was doing the blocking of the 5-dot address on those protocols, but not on ICMP. So it might be worth looking at. But that's what's going on is it's just - it's pure legacy. There's legacy tables in routers on the Internet that have not been updated to reflect the fact that ranges that used to never be seen are now being seen. It was really never necessary

for anyone to block those. But somebody was just a little bit overzealous, and they did.

Leo: There you have it. Our bonus from Mike in Daytona, Florida: There is a new Peter F. Hamilton book, the "Great North Road," 900 pages big. Have you already started it?

Steve: Oh, no. All my time is being sucked into catching up on the prior three seasons of "The Good Wife," which I'm enjoying. And when I'm not doing that, I'm plowed into the TrebleShooter project for the Quiet Canine effort, and again, making fabulous progress. I'll have a full update and get the pages updated before long. But I did go and look on Amazon. It's getting well reviewed. People are not all five stars. It is long. Apparently it's a mystery, which is interesting. It also involves procedural police work, which he had in the "Pandora Star" series, the two-volume. Remember he had Paula was this really cool...

Leo: Investigator, yeah.

Steve: ...investigator. And so there's some of that. And there's mystery. There's chasing aliens around. This is in the far future, where remember that in Peter's previous work, where we had - in "Pandora's Star" they had wormhole technology and were running trains through wormholes, which I loved. Here, there's still no faster-than-light travel, but they've figured out how to do teleportation. But the teleportation is controlled by one family which is multiple generations of clones. And the clones, as will happen if you keep cloning clones of clones of clones of clones...

Leo: Transcription errors.

Steve: Yes, you get exactly, DNA transcription errors. So they're becoming a little defective. Anyway, it sounds like, I mean, I just trust Peter to give us a really romping story.

Leo: But it's still got sci-fi, it's just a sci-fi detective story.

Steve: Oh, it's absolutely science fiction. It's set in the far future and in a really interesting universe. So I can't wait to read it. I haven't yet, but thank you, Mike, for bringing it to our attention. And some people also tweeted it to me, so I had seen that before I encountered it in the mailbag just now.

Leo: Audible, let me check, they have a lot of Peter F. Hamilton stuff. I don't know if they have that one yet.

Steve: If not, it's only because it's so new.

Leo: Yeah, and it's 900 pages. Nontrivial. "Great North Road," is that the name of it?

Steve: Yes.

Leo: Yes, they do. Unbelievable. Well, I guess I'll be listening to it, not reading it.

Steve: Do they say how many hours it is?

Leo: 36 hours. It's not that long. Not that long. I mean, that's long-ish.

Steve: Good. Please listen/read it for us, Leo, and let us know what you think.

Leo: Here's what the narrator sounds like.

[Excerpt]

Leo: I love Peter F. Hamilton.

Steve: Yeah, he's so good.

Leo: He's our favorite, one of our favorites.

Steve: He really - I would say he's absolutely, I mean, I love the work of Mike McCollum at Scifi-AZ.com, yeah. Love his work. But Peter's, oh, boy, I've got to catch up because I haven't done any of the Void Trilogy yet, either.

Leo: Oh. God, that goes on and on. Not my favorite Peter F. Hamilton, actually.

Steve: I know. Yeah. I think, if I were to recommend...

Leo: His weakest series.

Steve: I would recommend, if someone wanted to see if they like him, "Fallen Dragon."

Leo: Yeah, that's the one.

Steve: I've read it three times. It's just, oh, it's just...

Leo: That's a masterpiece, yeah. The Void series I would say not a masterpiece. "Pandora's Star"...

Steve: And then "Pandora's Star."

Leo: ...very good, very good.

Steve: Yup. And "Judas Unchained" is the sequel to that. And then you're pretty much done with blobs that become intelligent.

Leo: It's a great premise. I do enjoy the premise of that.

Steve: Oh, god. I love his world.

Leo: Yeah. Not so crazy about the Void Trilogy.

Steve: Okay.

Leo: Steve Gibson is at GRC.com, Gibson Research Corporation. See how that works out? That's where you'll find all his freebies and his bread and butter, SpinRite, the world's finest hard drive maintenance and recovery utility, even for solid-state drives. He also makes 16Kb versions of the show available there in audio and transcribed each week. Elaine transcribes every episode, and the transcriptions are online there. GRC.com. It's also where you'd go to ask a question, GRC.com/feedback. And there's lots of, I mean, it is a treasure trove of eclectic material that Steve has gathered over his eclectic life.

Steve: You and I, baby. Oh, you mean the site. Yes, it is getting ever more eclectic, actually.

Leo: Yes, that's what I'm saying. That's what I'm talking about. So if you like what you hear here, oh, you get a whole lot more there. In fact...

Steve: Yeah, I've often thought that it's nice that I called the company Gibson Research Corporation because it gives me complete freedom.

Leo: Right, it's whatever Gibson's researching this week.

Steve: Whether it's ketones in your breath, or it's the longest repeating strings, or it's how to quiet your neighbor's dog, we've got it all.

Leo: You get the idea, folks? Now you understand what I mean when I say "eclectic"? Okay. He also does this show, and we are very grateful to him for it, every Wednesday, 11:00 a.m. Pacific, 2:00 p.m. Eastern time, 1900 UTC on TWiT.tv. Watch live if you can because there's often stuff before and after that's worth staying tuned for. But if you can't, we make audio, as I said, audio and video versions available after the fact. Steve's got the baby ones and the transcriptions. I've got the higher quality audio and the video at TWiT.tv/sn. Or just go to iTunes or anywhere podcasts are offered, and you can search for Security Now! or TWiT. You'll find all our shows on all the usual suspects. Steve, I'll see you next Wednesday. Thanks so much.

Steve: Thanks, Leo.