Listener Feedback #162

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-393.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-393-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson's here. We've got questions from our audience. But before we get to those, I can't believe it, there's another flaw in Java. Details next on Security Now!.

Leo Laporte: It's time for Security Now! with Steve Gibson, Episode 393, recorded February 27th, 2013: Your questions, Steve's answers, #162.

It's time for Security Now!, the show that covers your security and privacy and everything else you need to know to stay safe online with this man right here. We call him the Explainer in Chief, Mr. Steve Gibson of GRC.com. Hello, Steverino.

Steve Gibson: Hey, Leo. Great to be with you. Episode 162, I'm sorry, 393. Q&A 162.

Leo: Question, yeah, Question & Answer 162, yeah.

Steve: Yeah, Q&A 162. We've got so much to talk about, sort of toward the top of the show, that I'm not sure how many questions we'll get to. We've got 10 plus a bonus question this week, some interesting, fun things to talk about. So we'll just do 90 minutes' worth of stuff, starting with news and updates and so forth, and play it by ear.

Leo: That's fantastic. That's fantastic. We have some visitors in the studio. Lyle's from Nashville, James from Bloomington, Indiana. I warned them. I said, "Are you ready to get geeky?" They put on their little pointy hats, and they're ready to go.

Actually, James is an IT guy at Indiana - is it Indiana University? Yeah. And Lyle's with Cisco? Works with Cisco. Did you see the small town that has a $30,000 Cisco router? Did you see that story? Just ridiculous.

Steve: No. So, like, so total overkill for what the town needed?

Leo: There's 400 people in the town. They have a shack for a library with one Internet connection, and it has a massive Cisco switch in there. And the reason is they went to Cisco, the state - I'm going to - got to find this story. The state went to Cisco, and Cisco said, yeah, you need all that stuff. They just threw it at them. It's really kind of an amazing story. I'll see if I can find it while we're talking. I just read it this morning, and I should have bookmarked it, fool that I am. But you've got plenty to talk about, so...

Steve: Well, you have an amazing story, and for the top of a podcast we have an unbelievable story, only because it is too believable.

Leo: Uh-oh. By the way, it's West Virginia. It's a small town in West - here, I just found it. It's in Ars Technica. One-room West Virginia library has a Cisco 3945 router. It's a $20,000 router. It's in a temporary building. It turns out, in fact, the state's legislative auditor found this. It's the small town of Clay, West Virginia has seven - seven - 3945 routers serving 491 people. $100,000 worth of Cisco routers within 0.44 miles of each other. This is the best routed city in America.

Steve: Oh, goodness.

Leo: Yeah, I guess it was Cisco's sales guys that convinced the state to buy $24 million worth of 3945 branch routers. Unbelievable. You couldn't even saturate one of them.

Steve: No. No.

Leo: So your story now. I'm sorry. I didn't mean to interrupt.

Steve: Believe it or not, Leo, we have two new Java vulnerabilities.

Leo: What? No, that's not possible. We fixed Java just the other day.

Steve: Yes, we did. And the day after, two new ones were found. A pair of newly detected flaws in Java can be exploited to allow attackers to bypass the browser plugin's sandbox security feature. So it gets out of the Java sandbox, which of course is now the goal of any exploit. It affects the most recent Java update, which is Java 7 Update 15, which was released on February 19th.

Leo: [Laughing]

Steve: So here we are...

Leo: Eight days later.

Steve: ...nine days later on the 27th. And actually this news is a couple days old. And Java 6 is not affected.

Leo: Oh, so it's something they introduced.

Steve: This I, well, yeah, exactly. This is something, and we've seen this already in Java 7, things that were not a problem in 6, new features they put in that are causing problems. So in several of the reports of this, I've seen it written that experts are advising users to disable or even uninstall Java. And there are also reports that an exploit for Java 7 Update 11 has been detected in the wild, and Update 13 was released to fix that one on February 1st. So it's really getting ridiculous.

And one of my favorite sources of security news is, you know, the SANS security newsletter that I subscribe to. And it's just a really great summary that comes out a couple times a week of stuff. One of their editors, John Pescatore, who was a VP of Gartner for 14 years, he had a little editorializing after their news summary. He said: "In developing the Morse code, Samuel Morse assigned the shortest code element (dit) to the most common letter (E) in the English language. Java vulnerabilities today, much like IIS" - and that was Microsoft's web server back in the day - "much like IIS and Internet Explorer vulnerabilities 10 years ago, have now earned 'dit' status." And I was tempted to say, "Or maybe 'duh.'" Just amazing. I mean, this is - it is getting to be ridiculous.

In fact, we've got - one of our listeners shares with us his company's solution to this problem, how they've worked - how because they have a need for Java, as corporations do - you know, invariably when I tweet something about this, people say, oh, well, you know, the best way to fix it is uninstall it. It's like, yes, but some people can't.

Leo: Well, say Minecraft, there's lots of reasons you need it.

Steve: Yeah, there are, actually. And in fact the code that I'm - the Eclipse platform, which is the development...

Leo: Is Java.

Steve: ...platform for - is Java based. So you've got to have it around. You don't always have to have it in your browser. Although, and we've heard reports from, like, lots of Scandinavian banks require Java plugins in the browser. So that's the problem. And, I mean, it's just - it is really getting pathetic and ridiculous. We do have a bright spot, however. Firefox v7 - v17 - version 22 - I think we're now - I think I'm on 18 or 19.

What am I on now? I updated not long ago. I'm on 19.0. Oh, and that's the one that's got the integrated - we talked about it last week, the integral PDF reader. Even though I had a PDF plugin - I was using Sumatra PDF and happy with it, although it kind of gets a little wonky sometimes - when I installed the latest version of Firefox last week, it sort of pushed the plugin PDF reader aside and took over with its own.

So, again, it's always good to be cautious of something new from a security standpoint because this is a lot of new code. But there's no way that anybody today could develop something which didn't have - develop something for the browser like the Mozilla guys that didn't have security as its No. 1 feature. I mean, they're just - I'm loving the fact that we're seeing the kind of inevitable escalation of awareness because that's the driving force behind these things being fixed in the long term. So that's all good. Anyway, my point is that with Firefox 22, which will be three versions from now, slated for a June 25th release, and I don't know how they know that, but June 25th...

Leo: Yeah, because I'm on 18 right now. I mean, I'm way behind that. And that's just updated.

Steve: On Safari.

Leo: Oh, Safari.

Steve: I'm sorry, Firefox.

Leo: I'm Firefox 19. And it says it's up to date on the Mac, yeah.

Steve: Yes, yes. So they are once again going to try to disable third-party cookies by default.

Leo: Good.

Steve: Yes. So they will join the singular rank - I don't know if you could have a rank if there's only one. But they will form a rank with Safari, the only two browsers then that have third-party cookies turned off by default. And I coined the term "the tyranny of the default" because it is so much the case that default settings are what everyone runs on. I mean, obviously not everyone, but the vast majority of people just assume that smarter people than they figured out how things should be set, and they don't want to mess with them because they don't understand them. And I've been tracking third-party cookie usage of GRC's visitors for years now, and I have a GRC.com/cookies/stats.htm. You can see a really graphic, well, because it's a chart, but a very clear display of the effect of Safari's decision on third-party cookie usage because, among all the browsers that I profile, one stands out as just ridiculously low level of third-party cookie usage because it's off unless people turn it on.

So, and this was Jonathan Mayer from Stanford. He's been mentioned in the podcast for years. He's a graduate student in computer science and law at Stanford, focusing on public policy, law, and the Internet and privacy. And he's got some interesting papers

written about the Do Not Track header and why it represents, when it's eventually implemented, the best of all worlds because...

Leo: Now, is third-party cookies and Do Not Track the same thing?

Steve: No. Third-party cookies are...

Leo: Because I don't see a setting in Firefox for third-party cookies at all anymore.

Steve: Oh, yeah. It's definitely there. And in fact...

Leo: I see "Tell websites I don't want to be tracked."

Steve: And that's good. That's Do Not Track. I think it's under the Privacy tab.

Leo: No, it's not.

Steve: Oh.

Leo: That's what I thought. I looked under Privacy and Security. I wonder if they've taken that out now.

Steve: Privacy.

Leo: Maybe they just turn it on - when they say "turn it on by default," maybe they mean it's on, it's on.

Steve: Okay. So on my Privacy tab I've got - oh, maybe it's because I've got "Firefox will use custom settings for history." And I think that maybe opened up a bunch of things. Because I've got...

Leo: Ah, use custom settings. Ah, now I see it. So you...

Steve: Bingo.

Leo: Okay. So the default, which is Remember History, doesn't have anything. But if you go to Use Custom Settings for History, then Accept Third-Party Cookies has a checkbox, and you can uncheck it.

Steve: Yup, yup.

Leo: That's different from Do Not Track, which was the - that was the issue with Internet Explorer.

Steve: Right. And it's funny because I was looking at browser headers yesterday with Chrome, and they're just like - there was DNT: 1. Just Chrome sending out its little beacon of, like, this user has requested not to be tracked. And it's like, oh, that's nice. So, yes, the Do Not Track is a statement from the browser saying that my user has asked not to be tracked. That's all it means. So it's controversial because...

Leo: It's not even enforced.

Steve: ...honoring it - right.

Leo: It's just a request.

Steve: Honoring the request is up to the people who receive it. Third-party cookies are...

Leo: Very famously, Apache does not honor the request; right?

Steve: Well, Apache actually goes further. Apache, and this is something I think is so wrong, the server itself strips that header preemptively so that applications on Apache can't even see if it was sent. And I think that is really overstepping the bounds. This will end up getting reversed because it is absolutely wrong. But, you know...

Leo: I do wish Firefox would make it easier to turn off Java, to disable the...

Steve: Yeah.

Leo: I mean, come on. Okay, that's good...

Steve: Although, I think if you look - where was it I saw it? I think if you look in Plugins...

Leo: Yes, that's what - both Chrome, yeah...

Steve: Tools, add-ons.

Leo: Tools, add-ons.

Steve: And then if you look, if you do...

Leo: Yeah, there it is.

Steve: ...plugins, then it, like, it shows it in red and says "Known to be vulnerable." Acrobat, I've got a couple old Acrobats in here, 831, 930, known to be vulnerable, disabled, disabled. So, I mean, it really does - it's doing a good job, and we're going to talk in a minute here about something that Google did for the NBC.com site when it got hacked last Thursday that was really nice for people, too. So anyway, so Firefox v22 makes a very nice step forward thanks to Jonathan Mayer at Stanford. I want to, you know, tip of the hat to him and the Mozilla guys for saying we're ready to do this now. They did, at one point, it was, like, version, I don't know, 4 maybe, when it was in beta, they had it turned on. But they backed out at the last minute before it went public and said, okay, no, we'll just leave it the way it is.

But we're seeing a groundswell about privacy. And I'm seeing some things from the people who really track this about how privacy and security are beginning to overlap, that is, people may say, oh, well, privacy, that would be nice to have, but no one has it. The problem is attacks are becoming more targeted. And the more information the targeter has about the targetee, the greater the chance that the attack is going to be successful. So it really does make sense to keep your privacy guards up where you can. So anyway, that's nice news for Firefox. I know that, I mean, I guess the people who are in the know today are probably either using Firefox or Chrome. Certainly no one's using IE, although it's too bad because Microsoft has really, you know, they've...

Leo: Hey, you know what? I've been using Windows 8 on a high-res tablet, and Chrome looks like crap. So I'm using Internet Explorer. Chrome is my default, but it does not look good on this 1920x1080 display, and it doesn't support touch. So I can't, you know, with Internet Explorer I can scroll with my fingers. I can zoom and pinch. And Chrome not yet. I'm sure it will. Maybe...

Steve: And the fact is, IE10 today, I mean, we're still...

Leo: It's good.

Steve: Yes. And Microsoft has really - they focused on security. It took them forever to get off the launch pad and into orbit, as always is the case with them or any really huge organization. But they have. And so it's really not fair to accuse them of still having IE6. They don't have IE6 anymore. They have IE10, and it's way better. So...

Leo: Yeah, I ended up using IE10 on Windows 8 just because it was better.

Steve: And speaking of which, IE10 just came out of beta. I upgraded to that on a

Windows 7 machine of mine yesterday. You know, 7 comes with 9, and I updated to 10 because I wanted to start watching it and see how it feels. It looks absolutely the same, so it's just - it's better.

Leo: I guess I can scroll with two fingers. It's just weird. You know, Explorer just looks better. Just handles it better. So if you're using Windows 8, I think there are probably a lot of people who've gone back to Explorer.

Steve: So what I was going to say was that on the desktop platforms we're probably split between Firefox and Chrome, either of which are really good choices, I think. What did not get any news coverage, I was really surprised, I didn't even - no one seemed to pick it up and tweet it to me, so my major pulse on the industry didn't see this, is that Microsoft was also hacked along with Facebook and Apple.

Leo: Yeah. They didn't mention it.

Steve: No.

Leo: Until everybody else did.

Steve: Exactly. And so in their blog Friday they said, under the title of "Recent Cyberattacks," they said: "As reported by Facebook and Apple, Microsoft can confirm that we also recently experienced a similar security intrusion. Consistent with our security response practices, we chose not to make a statement during the initial information-gathering process. During our investigation, we found a small number of computers, including some in our Mac business unit - in other words, it was the same iOS watering hole attack that got the Facebook and Apple developers - "that were infected by malicious software using techniques similar to those documented by other organizations. We have no evidence of customer data being affected, and our investigation is ongoing." Well, good luck with your investigation.

Leo: [Laughing] You're such a cynic.

Steve: I know [laughing]. I've been watching them for too long. Although I do, as I said last week, I've been very impressed with Server 2008 R2. And I am using a new Win7 installation with IIS 7.5 here at home so I have the same thing that I have at the data center. And I'm - they've got their act together finally.

Leo: Good.

Steve: Apparently you know someone named Bob Bosen, who you once talked about his AskMrWizard.com website.

Leo: Sounds right.

Steve: He had some interaction with you. Well, Bob asked me if it would be all right with me if, for some selected propeller-head episodes, which sort of cry out for graphics, if he were to do some animated graphics to go along with my voiceover.

Leo: Oh.

Steve: And I said of course.

Leo: Yeah, I love that idea.

Steve: So you can go to AskMrWizard.com, and over on the left you will see a link to Security Now!, or you can go AskMrWizard.com/securitynow, which takes you directly there. And what I saw yesterday, and I don't know how he's organizing it because I just saw his note when I was going through the...

Leo: He calls it "Security Now! Illustrated" [laughing]. I like it.

Steve: Yeah. And he handles NAT traversal. He talks about NAT routing and how NAT routing are good firewalls. Basically he's using the podcast audio, but doing 3D animated video to...

Leo: Wow.

Steve: With, like, comments and balloons that illustrate things. He understands it, and so he's helping other people to do so. So what I'm going to do is I want to tell everybody about it. I asked him, can we tell everybody? He says, oh, yeah, sure. So this stuff is up on YouTube. And what I would hope is that, as he finishes new ones, he'll let me know, and I'll just make a brief note during the show that there's another one for people to check out, if they want. But they're nice. I mean, it's done at home, but it's nice 3D graphics, animated stuff.

Leo: That's really cool. That's so neat. And I do agree. I've always wished we had the budget to have more illustrations in all of our shows.

Steve: Yeah. And especially for this, when we're doing the fancy stuff, when I'm talking about encryption and all that, it'd be wonderful. But, oh, my god, that's - that would take over my life for the week, getting that stuff done because...

Leo: Oh, yeah, yeah, yeah, you can't.

Steve: ...I'm a perfectionist, and I would never be happy with it. So, okay, now, Leo, this is beyond cool. You need to click the link in that next item.

Leo: Okay. Don't say it out loud until I do.

Steve: I won't.

Leo: Because we'll probably bring it down.

Steve: This is something that someone following me over in my VLC Twitter feed tweeted. And it is so cool. It's called the UChek Urinalysis.

Leo: [Laughing]

Steve: And it is brilliant. They use a smartphone, like an iPhone, to - you take a picture of urine test strips over time. So you take several different pictures of it as the test strip is evolving. And then this thing interprets the results for you much better than you could just by hand, and of course builds in a calendar and all that. And it measures ketones, all kinds of measures which are detectable from a urinalysis.

Leo: You know, this is nontrivial. I look at these things, and I don't - is that purple? Is it green? I don't know.

Steve: Yeah. And, now, okay. And the thing that I got the biggest kick out of was the developers who said, you know, this is going to be an Apple application, so we're tempted to call it iPee.

Leo: iPee. There will be an Android version, too, though, so iPee.

Steve: Anyway, just...

Leo: I like it.

Steve: Yeah, brilliant.

Leo: Not out yet, by the way. But UChek.in, if you want to check for it.

Steve: Ah, good, yes. Looks neat. Okay. So this is, again, we're doing miscellanea here. But this is the most exciting piece of news I've seen in years. It is the cover of next week's Time magazine. It's been on the Time.com site for about a week. It is a long story. I tweeted the link to it. I also created a bit.ly shortcut, bit.ly/bitter-pill. "Bitter Pill"

is the title of this Time magazine - it is a major investigative report, the longest one Time magazine has ever undertaken, seven and a half months they worked on this. And what's exciting to me, that whole title is "Bitter Pill: Why Medical Bills Are Killing Us." And this has been my own personal crusade, I mean, my friends are all annoyed with me because I'm ranting because, you know, here we are with the Republicans and the Democrats locked up, we've got the sequester coming into effect on Friday, and everyone's yelling and pointing fingers at each other. They're worrying about entitlement reform, and the entitlements are going to - the soaring cost of medicine. And I have been asking, okay, the driver of all of this is that it costs so much. Why does it cost so much? Why does, okay, yes, open heart surgery is amazing.

Leo: Yeah, but a Band-Aid shouldn't cost $25.

Steve: Exactly. And I have an elderly friend who lives in my neighborhood who's had a couple rather routine surgeries, but because he's 75 they kept him at our local Newport Beach Hospital.

Leo: Oh, you don't ever want that.

Steve: It was, like, $7800 a night. $7800 a day. And it's, I mean, just mindboggling. And it's like, okay, why? And what I'm excited about, and the reason I'm bothering our security-oriented listeners about this, is I don't think there's anything more important. I have no problem with Mercedes and BMWs costing whatever they want because I can choose if I want to drive a VW or a BMW.

Leo: It's discretionary, yeah.

Steve: Yes. But health isn't. And this article, which is fascinating, I'm mentioning it because I would encourage everyone to make some time to read it. It'll be out on the newsstands this weekend. It's next week's Time magazine cover story. You can get it now online at Time.com. And it is fascinating. This reporter doesn't just wave his arms around and talk about how expensive everything is. He goes through individual people's case histories, talks about their - dissects their 327-page medical bill and demonstrates how wrong it is that a hospital in Texas that is nonprofit, and therefore tax exempt, is making a billion dollars lobbying Congress. The medical industry lobbies Congress five times more than the defense industry does. It's amazing.

And so anyway, I just wanted to point people to it. Enough said. But that's the problem that we have is it's the expense at the far end that is pushing everything back. So I just thought it was really interesting. And I was so glad to see it because we have to have - this has to be brought to light. And nothing will do it better than a cover story in Time magazine. All the different news stories, Jon Stewart had the guy on late last week. All of the morning shows have been talking about it. So I'm just glad it's getting some attention. It's really important.

Also, one of my very favorite sci-fi authors is now on Audible, Michael McCollum of Scifi-AZ.com. So I wanted to let people know. I got the heads-up from a couple of our listeners who said, hey, Michael's books just appeared. And by "just" we're talking, like, the 19th, you know, so just last week. I pinged him and said, hey, what's the story?

Because there are five of his books are there now, the first two of the Antares Trilogy, the first one of the Gibraltar Trilogy, and both of the Makers series. But they're all going to be there. He told me that Audible had purchased the rights to all of them. He was removing his own Audible experiment where he was - I think he had an automated text reader that was, like, reading them, which is way less wonderful than an actual person reading them. But they're out there now. They're on Audible. And so I just wanted to give our listeners a heads-up. The Antares Dawn series is just great, as is Gibraltar Earth. I really, really enjoyed them.

Leo: I'm so glad they're on Audible. Did he make a - he must have made a deal because he's his own publisher; right?

Steve: Yeah, he is. They bought the rights. He has an agent. And so the agent made a deal with them. They wanted them. And so I'm delighted that there was the interest, yes.

Leo: And I presume they're going to do Book 3 of Antares. They've got to.

Steve: Oh, no, they're going to do them all. All of his books will be there. "Antares Dawn," the first one, was released on the 17th of this month, February 17th; "Antares Passage" on the 19th.

Leo: Oh, okay.

Steve: So, and then "Procyon's Promise" on the 20th. So they're all on their way out.

All the PDP-8 clones were sold, Leo. The 19...

Leo: Oh, shoot. I just missed out.

Steve: [Laughing] Yeah. The ultimate nothing box, as you would - I can't argue, but they're just wonderful. And many people have been writing, talking about something that's been in the news, but is really not that newsy. But so I wanted to cover it just so that people would understand sort of where we stand. And that's over graphene. Graphene is amazing. But it's about eight years ago graphene began to happen. What was newsy was that a video that was made of a bizarre way that a couple guys in the lab created graphene by, like, dripping it on a CD and putting it in their CD burner and, like, using...

Leo: What?

Steve: Yeah, they used the UV for the laser of...

Leo: Oh, that's hysterical. Homemade graphene.

Steve: ...to, like, create the graphene. And they produced - and there's a beautifully produced video that sort of doesn't get into as much detail as I would like. But it won some awards for, I don't know, great graphene documented something.

Leo: It ain't the Academy Award, but it'll do.

Steve: But what's exciting, and the intersection with us and me and the podcast, is there's probably nothing more perfect for supercapacitors.

Leo: Oh.

Steve: Now, what graphene is, it is an amazing substance. It is pure carbon, so it is just carbon molecules interlinked in a hexagonal matrix which is one atom thick.

Leo: Wow.

Steve: So imagine just a flat sheet of carbon atoms linked to each other absolutely flat. It's called an atomic monolayer. But get a load of its properties. Okay, in terms of resistivity, that is, its resistance to electrical flow, a graphene sheet is 10E-6 ohms per centimeter. So...

Leo: Very low. Very, very low.

Steve: Unbelievably low resistance. In fact...

Leo: Is it a superconductor in that case?

Steve: It's not, but it is, like, it's, like, close. It's less resistance than silver, which is the lowest resistivity substance known at room temperature, and graphene is lower than that. It's also one of the strongest materials known. Get this. A square meter graphene hammock, imagine that you made a hammock that was a square meter, so a little more than a yard by a yard. That would support a 4-kilogram cat, an 8.8-pound cat. Okay?

Leo: Okay.

Steve: It would weigh only as much as one of the cat's whiskers.

Leo: Wow.

Steve: 0.77 milligrams. It is that thin and that low weight, but that strong. So that a square meter of graphene sheeting could support an 8.8, almost a 9-pound cat, yet

weigh about as much as a whisker. And it is .001 percent of the weight of paper, .001 percent. So just incredible. It is also transparent. It only absorbs 2.3 percent of the light passing through it.

Leo: Oh, see, I was going to make a solar sail out of it, but it wouldn't be very good for that.

Steve: Well...

Leo: Strong but invisible.

Steve: Because you need to be reflective in order to...

Leo: Yeah, to capture it.

Steve: Yes. But think of all the instances where we want something conductive that's an electrode, like a touchscreen or an LCD panel, where we need to pass electricity. Nothing does this better than graphene. So huge applications there. But more than anything else, ultracapacitors. What you need for an ultracapacitor is two electrically conductive surfaces very, very close. And it turns out you can create graphene oxide, which by being an oxide is an insulator. So you take a graphene sheet, a graphene oxide and a graphene sheet. And they are incredibly close together. They're one molecule apart. And the graphene sheets are incredible low resistance. So they can hold a large charge with a high breakdown between them.

Initial studies of using graphene for ultracapacitors has shown that it has more energy density than current lithium metal - it wasn't metal - it wasn't lithium ion. It was lithium metal hydride, which is slightly lower energy density than lithium ion cells. But still, real, I mean, like, this is existing, I mean, we're getting this in the lab today. So this isn't, you know, bizarro future technology that isn't close. And everyone's talking about cell phones that you can charge in a second whose charge lasts a day. No more of this hours to recharge. Or cars where you drive them into a high-current charging station, plug in for 60 seconds, and you are completely topped up.

Leo: Yes, that's what we need.

Steve: So it is looking like graphene is the way we're going to get there. It is just an amazing substance. And apparently you just squirt it on a CD and spin it in your computer, and you're able to make some of it.

Leo: That blows me away. Not a usable amount of it.

Steve: Just, you know, it's a...

Leo: It's a graphene spot.

Steve: Yeah. I'm not impressed by that because it was fun in the lab, but in commercial settings they're not going to get a roomful of CD drives, keep burning.

Leo: It means it's not a very, very difficult thing. You can do it at room temperature with simple devices. So that's good news.

Steve: I mean, everybody is racing. The Patent Office is under siege with graphene-related patents. So there are a lot of things that have been done that no one can talk about yet because they're waiting to, I mean, they're having to go through some intellectual property protection process. But labs everywhere are going nuts. And it turns out that maybe it's going to affect both semiconductor production, because you need conductive surfaces on top of semiconductors. Right now they use what's called "metalization layers." And graphene, if they could figure out how to tame it, could be way better than what they have now for semiconductor metalization layers. And there's been some notion that this may create a breakthrough in quantum computing, as well. So this is really cool stuff.

Leo: Graphene. Now, Steve, are you ready? See, I knew you wanted to know what this box was.

Steve: That's the question I was waiting for, Leo.

Leo: Are you ready for questions?

Steve: Am I ready.

Leo: Here we go. Question 1, John Thompson, Albuquerque, New Mexico. Look, he spelled it exactly right. I guess that's because he lived there. He suggests that the canine isn't the only thing that's been quiet about the Quiet Canine project: Steve, it's been months since we've heard anything definitive about your work on the Quiet Canine project and your development of the TrebleShooter devices. At one point you said that hundreds of listeners had used the Quiet Canine feedback page to tell their stories and plead for a solution. I number among them. So perhaps you'll consider this as a short topic for a forthcoming Q&A? Where is the canine?

Steve: So, okay. I did get pulled off of it by the emergency that I had with my 13-year-old servers that began to get too long in the tooth, I guess. They started having weird problems that were of no discernible cause. I moved them to different hardware, same problem. I rebuilt the software, no improvement. I just think it was - actually I think it was date related, and it began on January 1st, and it kept getting worse until toward the end of January it became a real problem. So as I have mentioned before, I did have brand new servers waiting to be deployed, and this was finally the impetus to make that happen. So I have a few things I need to clean up to finish resolving. There were a bunch

of incompatibilities that I feared. All the major ones were cured quickly, and I've still got some I need to clean up. Then I'm going to get back to just wrapping up that project.

But here's what I think we've learned. And this is from all the experiments that people have done so far, even though there's still one thing left I really want to try, and that is to try a tuned resonant cavity with a transducer, a piezo bender, or a tweeter driver at the end where we tune the cavity for resonance and see how that works. To do that, we need a high-voltage variable frequency sine wave source, which is what I was working on when I got pulled off to fix the servers. But what I think we know now is that what many people want is not possible.

Leo: Oh, no.

Steve: That's the rip-the-bandage-off reality. What people want is something that will silence a dog three houses away or two stories below them in an apartment building. And we don't have that.

Leo: Yeah, okay.

Steve: And we're not going to have that.

Leo: There's no sniper version of the TrebleShooter.

Steve: Well, and I was wondering why so many people were asking for that. And then I realized it's because several different aspects of the original Portable Dog Killer story kind of fit together that way, but they weren't really meant to. When I caused the rabid German shepherd attack canine to have its legs collapse out from under it and make an amazing sound, and then over the course of about a week and a half trained it never to rush the fence, it was by shooting it point blank in the face. I mean, I was at the fence. It was leaping over the fence, and I shot it right in the muzzle and elicited the reaction that I had. So it was - and as much as anything it was terror. I mean, this scared the dog, didn't hurt it, just scared the bejeezus out of it. And so that was the key, something it was absolutely not expecting really, literally, right in its face. And then the second adventure was with the seagulls, where we discovered that at a good distance away we were able to alter the flight path.

Leo: Hey, somebody just bought a copy of SpinRite.

Steve: I actually have the phone next to me. It now works out of the BlackBerry as well as in the living room. So wherever I am I get a yabba-dabba in my pocket. Quite nice. Anyway, so there was the altering the flight path of the seagulls. The problem is, sort of by combining those two effects, people have been hoping that they could stop a dog from barking at the same distance that the seagulls had their flight path altered. And it's not the case. The dog might even bark more because it's going to hear this and think, huh, what's that? I think maybe I'll bark at it. So, unfortunately, we don't have an answer for that.

Now, the good news is, many of the people who did send feedback through the Quiet Canine feedback page have a need for what I would call "tactical personal defense." And I think we have an answer there.

Leo: Oh, boy.

Steve: So, and it costs about $8. I mean, it's amazingly inexpensive. And that's just the little seven- or eight-component design that I have and a $2 tweeter with a couple, I think it's a 9-volt battery it runs on. And it's fantastic. So, I mean, and that would be if you're jogging and there's a vicious dog approaching you, that sort of thing, personal tactical defense, where you need to discourage - I mean, and it's surprising how many people have this problem. So that we have.

Leo: Like a mail carrier would need this.

Steve: Yes. In fact, we have mail carrier listeners who have sent feedback saying, Steve, I need this. I'd be happy to test it out and tell you how it works.

Leo: Because right now they carry some - Mace or something to get the dogs that...

Steve: Yeah, and, you know, that's horrible. That's...

Leo: Yeah, that's really mean, yeah.

Steve: Yeah. Now, the alter- the other thing - so that works. The idea of a handheld personal short-range tactical defense, that we have because that's essentially the same thing that worked back when I was 19 with the dog that was attacking me. And it just stopped it cold. And I have a feeling this would work, be incredibly effective there. But the problem is, at a distance, there's just no way to project something really, really powerful that isn't going to fall off as it carries. I mean, the dog will hear it, no doubt about it. But it's not going to stop it. Now, we do have a...

Leo: It might anger it, and that wouldn't be good.

Steve: It very well could. We do have some dog owners who want to stop their own dog from barking. And there I am - what I'm thinking of is a dog collar. So you get the advantage of proximity and the electronics module sitting down by the dog's larynx so it's able to discriminate the dog's bark from alternative dog bark, other dogs barking or other sounds and so forth, and it would have a microcontroller in it and be smart so that it's not going to misfire. But then there are some, sort of some surface mount tweeters. Pyle makes some very inexpensive little surface mount tweeters, two of them. I guess people, like, stick them on their dashboard in their cars to get extra high-end response from their stereos. But if you stuck those on either side of a collar, then that would also definitely deter a dog from barking.

The problem, of course, is it wouldn't deter a neighbor's dog from barking unless you could say, do you mind if I put this anti-barking dog collar on your dog? And if they said yeah, we don't care, then that would work. But the idea that you're going to shoot something three doors down and stop a barking dog, it's just - we don't have that. That's just not going to be possible.

So that's where I am. I expect a few weeks from now I will have all of the GRC debris wrapped up, and I'm going to then get back and wrap the project up, make the pages public, make the designs public. Probably still going to build a bunch to provide to our listeners who have an application that would fit this. So that's where we are.

Leo: Excellent. Question 2 from Matt, who is at a secret location. He offers his Java strategy. Steve, I thought I'd chime in on how my organization is handling Java. I think it might be helpful for others in the Security Now! audience. I work at a research institute housed within a big university. We recommend that all of our users use Firefox or Chrome with Java plugins disabled for their primary browser. The university has some core services that require IE, Java, and in some cases both. So we leave Java enabled in Internet Explorer and ask users to use IE only to interact with those particular services. This way our users can still get their jobs done and are better protected during their other regular activities. And I guess because they're using Intranet sites they're probably not anything to worry about using IE in that case. Thanks for Security Now!, SpinRite, and everything else you do. That's probably a good solution, if you can get people to do it.

Steve: I liked that, yeah. I mean, all Windows machines have IE in them, whether you want them or not, thanks to the weird architecture that Microsoft created in order to say that, back when the DoJ was fighting them for antitrust, and they were trying to say, oh, no, no, we can't remove Internet Explorer, it's part of the operating system. It's like, what? So they made it part of the operating system for no good reason. And, yes, Leo, I am a cynic. But since you're going to have IE anyway, and you'd really rather not surf the Internet on IE, why not put Java there, like bundle all your problems in one place and then not use them at all.

Leo: Yeah, only use it on the Intranet or something like that, yeah.

Steve: Yes. Another thing, your comment about that read my mind because IE also has that whole notion of security zones.

Leo: You could say only on trusted sites.

Steve: Yes. So you could crank up the Internet security all the way to high, and there's that new feature that they have, that enhanced security thing that prevents you from doing anything useful. You can't do anything if that's turned on. But turn it on, because you don't want to do anything by mistake, for the Internet. And then for Intranet, where it's a block of IPs that are local, hey, then you're able to use Java in IE with no trouble. So I think there's a way to tame this, essentially, by giving Java its own playground. But just make sure that it's removed from the browser that you normally choose to use.

Leo: Yeah. So you have a local Intranet zone. And you could say you can only do stuff there; right?

Steve: Yes.

Leo: And when it comes to the Internet, turn it all the way up, and then you're all right.

Steve: I think that that must be a corporate, I mean, I'm wondering...

Leo: Oh, it's definitely for this purpose exactly.

Steve: ...how many users get into their zone configuration. And what happens is, using Active Directory and Group Policies, which are Microsoft technologies, the IT department is able to propagate settings out through the entire network and, like, remotely lock down and configure their browsers for their users. So it's good for that.

Leo: Precisely. Question 3 from a ham in Ireland, in Galway, EI8DRB: Steve and Leo, huge fan of the show. In fact, it was Security Now! that introduced me to the world of podcasting in general. But we like to call them "netcasts." But thanks for providing me with the company in the frequent traveling I do as part of my work. In recent episodes, Steve, you were talking about how parallelism has reduced the "hardness" of certain computational functions. You mentioned specifically iterative hashing functions and proceeded to describe how pipelining can reduce a multi-iteration hashing function to a series of discrete stages in a pipeline.

However, did you forget to take into account that each iteration takes the previous iteration as an input? Unrolling this into a multistage pipeline would be infeasible since stage N would require as an input the output of stage N+1, which in a pipeline would execute at the same time and thus would not have the result in time for stage N. Am I correct? Is there something fundamental I have missed? Thanks to you and all of you at TWiT and Security Now!. Regards and 73, Gerry. You talked about this recently, I think.

Steve: Yeah. I did. And I just wanted to make sure that this wasn't a common misunderstanding. Gerry is right, except that he got his math a little bit wrong.

Leo: N-1.

Steve: Yes, exactly. He said "since stage N would require as input the output of stage N+1." It's actually stage N+1 requires as input the output of stage N. So it works exactly correctly. That is, and this is in general the way pipelining of all forms works, whether it's CPU instruction process pipelining or any kind of loop unrolling pipelining. It is the output of the previous stage feeds into the input of the next stage, as opposed to where Gerry got himself tangled up. He had the output of the succeeding stage being needed for the

input to the preceding stage, which is not the case.

So it is the case that the pipelining works just beautifully. And it's just - it's such a cool concept that you're able to take something that is iterative and, as they say, well, and in fact in computer science terms it's called "unrolling the loop." One of the things when you are, I mean, in any kind of computer processing you often have a loop that you're doing multiple times. If you absolutely have to have the highest performance possible, there is in programming a slight overhead to the loop itself. When you get to the bottom, you typically have to check something to see if you need to do it again, like you have a counter.

So, like, say that you had to do something five times. So you decrement the counter. Then you check to see if it went to zero; and, if not, you jump back up to the top of the loop. So there's a decrement, a test, and a conditional jump which you execute every time through the loop. But if you really care about performance, instead you do what's called "unrolling the loop," where you actually have those set of instructions linearly laid out five times. Then you don't have to have a counter, you don't have to have a test, and you don't have to have a jump. The end of one set just falls right into the next set, which falls into the next set. And by the time it comes out, it's done. So you save the loop overhead. And that's, again, another form, sort of a form of pipelining relative to iteration, where the results, the work from each stage is sort of, by implication, is available to the next stage.

Leo: Question 4 comes from Ben in Whittier, California, just down the road from Steve. He wonders about RAID controller failure: Steve, I was listening to Security Now! 391 in which you talked about your monster server setup and all the redundancy and RAID 6 and all that. In my own admittedly limited experience with servers, the first component to die is usually the RAID controller itself. Are you increasing your exposure to such a failure by going with such an enormous onboard RAID controller cache? How messed up will your disks be if the cache memory goes out while the server is running? Thanks for explaining your setup. Best regards, Ben.

Steve: So, first of all, that's a very good point.

Leo: That is a failure point. I don't know if they're, I mean, they're solid-state, aren't they? They're less...

Steve: They're solid-state, and several things are going on. The thing that I am least comfortable with is that I have chemistry involved in the form of a battery. That is, you know, it's a lithium ion battery. It will keep the RAM contents in the RAID alive for two days. I really question whether I need it at all because I'm at Level 3 that has - when they were giving me their initial tour of the datacenter, oh, my god, they have a warehouse of batteries, just walls, I mean, just ranks of batteries. It's amazing. And then outside they have what looks like huge RVs which are industrial-size generators and tanks of fuel. I mean, these guys are really committed to the power never failing. So I really doubt that I need battery backup at all.

But maybe I would pull the cable, the power cord myself. Or maybe both of the redundant power supplies could fail at the same time. I mean, it's really unlikely, but you never know. So because I want the benefit of the so-called write-back caching, where nothing is written out to the SSDs until the space in the cache is needed, battery backup

is really the only way to do that safely. What the system does is every month it cycles the battery. I mean, it's very nicely done. It suspends the use of write-back caching and switches to write-through caching. Then, while it's doing this, then it takes the battery offline, drains it, and recharges it, measuring battery parameters all the while. And then, once it's back up to snuff, it returns to write-back caching. So it's monitoring the battery. It sends email reports of anything it finds.

But the worst could happen, and the RAID controller could fail. Which is why I have an entire second machine sitting idle. I bought two of the entire setup. So I've got - this thing's got redundant power supplies. I already have two spare power supplies in addition to an entire second machine, identically configured, sitting idle. I did that 13 years ago and never needed - I actually have three. One's running UNIX, one was running Windows, and a third one was just off, ready in case of an emergency. My feeling has been, since I would hate to be down for days, which I would be taken down for days, running around trying to find some other machine to run GRC, it's worth it to me just to invest upfront. So I have 100 percent redundancy of a system that is already ridiculously redundant.

Leo: If there ever be such a thing. Question 5, Jim in San Diego wonders how something like a dumb PDF viewer could be so dangerous. I can't - still, after our conversation last week, I look, and everywhere, everywhere says you need Adobe Reader. You need Adobe Reader. No, you don't. Oh, no, you don't. Steve, I love your podcast, listen as often as I can. As always, it seems, Adobe has fixed yet more zero-day flaws in Adobe Reader last month. Actually this month. Perhaps I'm not getting something obvious. Why can't Adobe Reader be written so that, even if it crashes, it would not have the ability to install malicious code? Is Adobe Reader's problem related to having access to kernel space? What's the deal? Jim in San Diego.

Steve: So, okay. And just to add to your comment about it being installed, I was setting up this new machine. I have a - it's a clone of the Intel Core i7 that I created when we were experimenting with that other video technology. It's a shuttle, nice little shuttle box. And I added the chips and RAM and so forth. And I was setting it up from scratch a couple of days ago, and I ran the installer for all of its own drivers, the chipset driver, the NIC driver, the display driver and so forth. And, bloop, there appeared on the screen a little icon for Adobe 9, you know, in that ruby red color of theirs. And it's like, yeah, you just cannot get away from it. Of course I uninstalled it and put Foxit Reader in, instead.

So the only way I can explain what's going on with Acrobat, or Reader, is that Adobe has such a stranglehold, for reasons that Leo and I were just saying, Jim, that it doesn't matter to them. Sure, some portion of the security-aware world is upset and doesn't use Reader, uses something else. But...

Leo: Everything has problems, though; right? I mean, it's not like there's a perfect solution.

Steve: Well, yes.

Leo: Software is software. It's just that Reader is ubiquitous.

Steve: Yes. And...

Leo: In fact, I might even say that one of the reasons you keep seeing Java exploits is because it's a great vector; right? So people are looking at it.

Steve: Well, and also it's that they are both incredibly complicated.

Leo: Right, right.

Steve: Complexity is the demon of security. I mean, it is...

Leo: And Microsoft has been aggressive about patching flaws on Windows. So you look to a third-party which is not updated so easily.

Steve: Well, and also what's the economic model? That is, Microsoft was clearly going to suffer with competition from alternative operating system platforms. Once upon a time there was BOS, and there was Apple and Linux and so forth. It was clear that Microsoft was suffering from their horrible reputation about security. And so they did the trusted security platform or something, whatever that was that Gates did his keynote on at Comdex [Trusted Platform Module].

Leo: Trusted Computing Platform, TCP.

Steve: That's right. And, but, okay, but Reader is given away. I mean, Reader is a free download. Java just, you know, we don't ever pay for it or ask for it. It's just there. So these are sort of weird loss leader things for these companies. I mean, Mark Thompson has been experimenting with the amount of money you can make if you bundle one of these default install checkbox things, you know, the Ask Toolbar, or what was it, the other thing? It was Ask Toolbar, and there was something else that Adobe was trying to install. Anyway, the point is...

Leo: Or McAfee. There's always something.

Steve: Yeah, McAfee, that's what it was. It's amazing how much money you can make.

Leo: Really? Good bucks, huh?

Steve: Oh, my goodness.

Leo: Mark Thompson is AnalogX. He does shareware. And he usually just gives it away.

Steve: Yeah. But he's like, hey, let's - I don't know. Something somehow crossed his radar. He did some research. And it's like, oh, my god, how much money you can make. And so here's Adobe saying, eh, we own PDFs. We invented it. It did go open source and open spec, so it's no longer a closed platform. It's getting installed everywhere, and we're making a bunch of money because we try to install crapware that people don't want at the same time. And it's like, what's their motivation for fixing it? They really - they're going to respond. But it's not clear that it hurts them. And the same thing for Oracle and Java. And Java is installing crap, too.

Leo: There are lots of third-party choices. People in the chatroom are wondering if there's one you like or recommend.

Steve: I do like Foxit. It's what I've been using recently. I installed it on Sue's machine when I rebuilt her computer a couple weeks ago and told her, okay, the icon changed now, but this is what it looks like.

Leo: Still use it.

Steve: Yeah.

Leo: I mean, it's had problems. It's had flaws, too. It's not free from flaws. But...

Steve: Well, that's just it. I would imagine, see, okay. One of the problems with Reader is it is ridiculously over-complex. Who needs JavaScript interpretation? I mean, talk about asking for trouble. If you're going to run a JavaScript interpreter in a PDF document, it's like, oh, just shoot yourself now. Because, I mean, we're trying to protect ourselves from JavaScript by locking down our browsers. And notice that the browsers, where their reputation and security matters, they fix it. Whatever it takes, they fix it. I don't see anything like that happening from Adobe and Oracle with Reader and Java. They just don't have the need.

Leo: Does it use Ring 0? Does it access the kernel? I mean, is that part...

Steve: No.

Leo: No. I don't think it does.

Steve: No. The problem is that it's - and that's the other part is that we're still in an operating system where users want freedom. And look at how developers chafe at the restrictions that iOS puts on them. And even users, it's like, hey, why can't I do this? Why can't I have LastPass running on my Safari browser in iOS? It's like, nope, can't do that. And so that's a restriction that annoys people, but it's where security comes from. Well, we don't have that, really, to speak of on our desktop platforms because they're open, and people want to be able to run anything they want on their computer. If that's what you ask for, then unfortunately it means you've got a - your applications are

empowered to do whatever they want to on the system, which means that an aberrant application, malware in whatever shape it takes, is similarly empowered to do whatever it wants because the system is designed that way. We want it that way.

Leo: We asked for it, yeah.

Steve: Until it goes bad.

Leo: Right. More power. More power. Question 6, an anonymous listener - you suggest the name Bruce, so we'll use that.

Steve: Call him Bruce, yeah. He said, "I'm anonymous."

Leo: Call me Bruce. Posed a few good questions about Bluetooth keyboards. I was using a Bluetooth keyboard out at a restaurant the other day, and this question popped into my mind. Is what I'm typing encrypted? Could anybody just listen in on my Bluetooth keystrokes? And if my connection's encrypted, how secure is it? It seems as if there could be a problem with some sort of statistical attack on the data that's going up to my mobile device. If the attacker knows the data is all keystrokes, perhaps they could make an educated guess about the frequency of which keys were being depressed, extrapolate out from there. How do they avoid that problem, if they do? Do they pad out the packets with lots of junk noise or something?

I know I'm not the most well-versed in the operation of encryption, so maybe this is a stupid question. But I'm curious how you clever security folks address this potential problem with a piece of what is for me everyday technology. Thanks, if you do answer my question. I love the show and undersign the rest of the blah blah blah. Keep up the great work and enjoy your new server hardware, Steve. You sound like a kid in a candy store, and your enthusiasm is infectious. Bruce.

Steve: Okay, so a couple things...

Leo: We talked about this some years ago.

Steve: Yes, we did. And I thought we hadn't for a long time.

Leo: Yeah, I'd love an update.

Steve: So I'll just dip a little into it. So let's talk about wireless keyboards in general first. The very early wireless keyboards had ridiculous, quote, "encryption," unquote. And you'll remember that all they did was, as you pressed keys, they sent out an 8-bit byte, which is the code for the key. But they said, oh, we don't want that to be easily readable, so we're going to XOR it with a byte. Now, we showed the graphic image last week or the week before, I think it was last week, where you XORed an image with noise, and you got noise. That's really good encryption. The problem was the keyboard XORed every

byte with the same byte. That is, it never changed. So every time you pressed an "E," you got the same code. It wasn't the code for "E," but it was a fixed code. And "F" was exactly one larger than "E" because it's ASCII. And so anyone who analyzed that stream, I mean, it'd be almost a nice, really amateur crypto problem for people to have. It's like, here's someone typing something, and we've XORed all of the bytes with a single fixed byte.

Anyway, obviously weak crypto. Anybody capturing that could easily, by doing a frequency analysis, look at which characters are typed, which characters are in relation to each other because it's not like there's any kind of cipher where each character maps to a different byte. It's that each character is just XORed. Which means the same - and remember XOR is a certain set of their bits are inverted. Always the same ones, always the same way, every time. So that's the way it was.

Now, but because that came to light, the manufacturers of the keyboards spent an extra three quarters of half a cent and increased the security of their keyboard so that it was good, so that it's now fancier security, and we don't need to worry about it. Like when you buy a Microsoft wireless keyboard or Logitech wireless keyboard, they've got good wireless security. But none of that is Bluetooth.

Bluetooth is a big set of standards, and the answer to the anonymous person whom we're going to call Bruce's question is that only in the instant of pairing of the Bluetooth keyboard with whatever device was receiving the keystrokes - I have a Bluetooth keyboard with my iPad, so maybe - and I love that little, the Logitech, it's like it forms both a cover and a stand, beautiful little keyboard for the iPad. And it's Bluetooth. Only during the brief event of pairing is there any danger at all from a man-in-the-middle attack, and even then it requires some sophistication. So what we said years ago when we covered in depth Bluetooth security, we've got a podcast on that if anyone wants to get this in more detail [SN-280], is Bluetooth has a known limited range of 30 meters. Wait, no, 10 meters, 30 feet.

Leo: 30 feet, yeah.

Steve: Right. Go out into the middle of an empty parking lot...

Leo: And you're safe.

Steve: ...where you can see all around you and you know that you don't have anybody within 30 feet or 10 meters. Pair your devices, and then you're fine. And also turn off discovery. Many people, it's shocking how many people leave their Bluetooth discovery turned on. I'm pleased that my BlackBerry that will soon be retired for a BlackBerry 10, when you turn discovery on, it's always for a fixed time. It just turns itself off.

Leo: Yeah, it should turn itself off.

Steve: Yeah, yeah. So the answer is Bluetooth keyboards are secure. Only if you were really concerned would you need to even worry about that pairing event. But presumably you do that at home or somewhere. And then, from then on, they know each other. When they are paired, what that means is that there is a secret handshake that they had

that established a shared key which is never exposed. Only derivatives of that key are ever used in subsequent communication, and you are as secure as we know how to make you today.

Leo: Question 7, Joe in Georgia - ah, I was wondering when we'd get one of these, a Universal Plug & Play question: Steve, I've enjoyed your discussions about UPnP and the external access issue. Quick question: Can this be solved simply by creating a DMZ pointing to a nonexistent IP? I seem to remember this was your advice when routers used to respond to IDENT queries, but I haven't heard anyone mention this workaround. Also, is there any risk creating a DMZ to a nonexistent IP on my internal network? I don't even understand why this would work. Is there any risk of a malicious sender being able to do damage to other machines, given that I've let them into the private side of the network? Really? Creating a fake DMZ, in effect.

Steve: Yeah. Actually, that's been a solution for a number of router problems in the past.

Leo: Because unknown traffic will just get routed to there.

Steve: Exactly. People have had routers where they were not stealth. One of the things that I did when I rewrote ShieldsUP!, as I talked about a few weeks ago, was I created that notion of true stealth, where I have a big funnel, and while I'm working with the user, sending them probing packets of different kinds, I'm casting a net for anything that comes back as a result of the queries that I am making. And I'd let them know, like, what's the nature of what bounced back for any reason. And so there were people, for example, whose routers did not allow them to turn off ping, and so that upset some people.

Anyway, so an interesting solution which arose is to create a DMZ, which is the - it's a funky, I don't know why it's called, I mean, what, for demilitarized zone. I'm not sure how it determines...

Leo: Yeah. Remember during Vietnam, there was like the place between your front line and the enemy's front line was the demilitarized zone? It was in between...

Steve: Actually, it's the Romulan Neutral Zone, Leo.

Leo: I'm sorry. Pardon me. Why don't we call it the RMN?

Steve: I think we should.

Leo: Or RMZ.

Steve: Romulan Neutral Zone. Then we would all know what we're talking about. So the idea is that anything not bound for a known computer behind the network gets sent to an IP that you establish. And you just set it to something impossible, something illegal,

something not in your normal DHCP auto-assigned address range, and it just kind of creates a black hole, which is a great solution. So we have encountered people who are unable to turn off Universal Plug & Play.

The reason I wanted, I had this question is in case any of our listeners have routers where they were unable to disable Universal Plug & Play as detected by, for example, my ShieldsUP! system. We're up to now 2,641 individual IPs have been found to be vulnerable. And the rate is really slowing down, which means that the listeners within the GRC sphere, largely this podcast and friends of friends of the podcast, have checked themselves and know one way or the other what their status is. But we're not seeing lots more. I mean, the number is apparently in the millions. But we don't have millions of people coming by.

So anyway, yes. The DMZ, setting up a router DMZ is a great solution for solving the Universal Plug & Play problem, routing port 1900 to a nonexistent machine. And then GRC will show that there's no response from you.

Leo: Good.

Steve: Great idea.

Leo: Clever. Jon in Sweden suggests a good answer for the anonymous email question in 391, a couple episodes back. Hi, Steve and Leo. Insert gushing praise here. In Episode 391 a listener asked about a free and anonymous email service. I think there's a great answer: TorMail.org is a free and anonymous service that cannot give your information away because they don't know anything about you except for what you tell them. And since their servers are off somewhere in onion land - that's what TOR stands for, The Onion Router - they are subpoena-safe. Also the problem Leo mentioned with exit nodes doesn't apply since it's a hidden service, and you never exit onion land. Oh, that's interesting. What a good idea.

Steve: Yes. One of the things on my list of probably propellerhead episodes is to describe the technology that TOR developed for offering services. Now, that's different than what TOR normally does. TOR is normally a client anonymizer, where client users use TOR to hide their identity for their traffic that exits out onto the public Internet to prevent them from being back-traced back to home. But the clever people who do TOR, I don't know if it was the EFF or whom, I'll find out and make the whole thing clearer in a podcast about it, they figured out a way to do the same thing for services so that you can access a service through the layers of onion routing, get to a server somewhere, but no one knows where it is. So it hides the service, and you are hidden from the service. So it's mutual hiding. Very cool stuff.

So for anyone who's interested in exploring anonymous email - and remember, TOR is not super efficient because you're running through, you're bouncing through lots of servers. People complain about the throughput, but I think that's because they're trying to download movies and things.

Leo: Yeah, but email it'd be fine for.

Steve: Email probably works very well. So Tor Mail, TorMail.org, check it out. I will, I will check it out, and we'll do an episode on how it's possible, what the technology is for offering hidden services, since that's very cool and not at all obvious how it would work.

Leo: Yeah. Matt Buford in Austin, Texas properly corrects Steve about the infeasibility of ISPs blocking Port 1900: I'm a bit behind in watching Security Now!. I just came across you suggesting ISPs should immediately block UDP port 1900. Was that for the UPnP problem?

Steve: Yup.

Leo: Yup. Whoops. Please don't suggest that packet filters be put in place for ports above 1024. These are commonly used as ephemeral ports, and you will block legitimate traffic by doing this. Of course, that's how FTP works; right? It just assigns a random port.

Steve: Yup. And that's how, like, client - well, I'll let him finish, and then I'll...

Leo: Imagine my machine wants to send a UDP packet to someone out on the Internet, DNS or whatever you want to imagine. It will pick up a random port at 1025 or higher, maybe 1900, as a source port. Many operating systems use 1025-5000. Sometimes people enable the entire range, which is of course 1025-65535 range. That's if they need a high rate of outgoing connections, servers, for example. Anyway, I'm able to send my packet out with a source of 1900 and a destination of something else, say 53. The response then comes back to me with a source of 53 and a destination of 1900, and your filter blocks it.

Steve: Yup.

Leo: If my OS uses 1025-5000 as ephemeral ports and my ISP blocks packets to me on port 1900, then one out of every 3975 requests of mine will fail, since I won't be able to get the response. I've found this problem on large networks more than once. People try to block worm ports or other things they think is bad on the Internet backbone and do so by just using ACLs to drop packets - Access Control Lists - to drop packets destined for that high port. Doing this causes random connections in the other direction to fail. The number of failed connections is often low enough to not be very obvious to the average user. But having an Internet connection that fails on even one out of 65535 connections would be a bad idea. Is it any better because it's UDP, Steve?

Steve: No. He's completely right. His example is spot on. He used the example of DNS, which is a UDP query. It often is the case, actually, that DNS is sourced from 53 and aimed at 53. It's often that 53 is to 53 port to port, but doesn't have to be. But his point is well taken. I was aware of this weeks ago and just kept forgetting to mention it on the podcast. So when I saw Matt's posting, I thought, yep, let's make sure we acknowledge that really it is not safe to block, and for exactly the reasons he says, anything over 1024.

ISPs, as we have said, are blocking, for example, 25 to keep their own customers from being spam generators. They are blocking - and that's outgoing 25. They're also blocking incoming packets destined to 135, 136, and 137, which are the old Microsoft NetBIOS ports. They're also often blocking 445, which is the newer version of Microsoft's file and printer sharing service, to prevent those old problems. That's not a problem because those are so-called in the service port range, below 1024.

For example, when you go to ShieldsUP! and you click on "All Service Ports," that's the block from zero, and I actually do check for port zero even though it's illegal, up to 1024. That's the block that I check because those are where services run. And so exactly as Matt says, when you first turn on your operating system, and it starts sending out queries, if you look at Netstat, which monitors the history of connections, you'll see 1025, 1026, 1027, 1028. It just marches right up using sequentially upward numbering ports until it gets typically to 5000. There's a registry setting you can change to change the upper limit, but that's normally enough. Then it wraps back around.

The reason that it works that way is that that disambiguates incoming traffic. You've got traffic coming back that may be TCP, and it may be for your IP, but the question is, for which query? Or, like, what is it coming back for? The only remaining way of separating many outstanding queries and replies is the port. So each successive query is on a succeeding numbered port. And exactly as Matt says, it's entirely foreseeable that you'd send something out on port 1900. So it would go out onto the Internet. And when the query came back, it would be coming back to your port 1900, which is exactly the port that Universal Plug & Play accesses on your router. And if the ISP were blocking it, you'd just get no connection.

Now, the operating system would retry, and it actually would use 1901. It would go to the next port and try it again. So our systems are robust around that, but it's still uncool. And he's absolutely right. The idea of this being done out on the Internet backbone just gives me shivers. That should never happen.

Leo: Uncool, man.

Steve: Yeah.

Leo: Uncool. Question 10: James S. in Utah wonders if his shields are only somewhat up: Thanks for the great Security Now! podcast. I've listened to every episode at least once. I just replaced my router with an Apple AirPort Express, since the old Netgear is now in pieces in the trash. I wonder if he hit it?

Steve: I think he probably hit it with a hammer.

Leo: With a hammer. Once I set up the new router I navigated to ShieldsUP! and started pounding on the thing, boom, boom, boom. The router responded to ping - fail - and actively rejected UPnP probes, but everything else was stealth. It turns out Apple will not allow us to turn off ICMP response on the WAN side of the AirPort or Time Capsule. Is that safe? Do I need to take my shiny new router back to the Apple Store? I didn't know this.

Steve: Yeah. This is a longstanding complaint.

Leo: Other routers do that, too, right, because...

Steve: It's varied. I think I'm responsible, actually, for the more standard Netgear, Linksys and so on routers disabling their ping response because of ShieldsUP!, frankly, and the idea that - and people just wanting to be stealth. So first of all...

Leo: "Stealth" means if a guy sends you a request, there's just no answer. Not even a no, I don't - I block that port. Just nothing.

Steve: Yes. And of course I named it "stealth" because of Star Trek and as related to the Romulan Neutral Zone and so forth.

Leo: Well, there you go.

Steve: So, okay. I don't want to alarm anybody. So, James, you're not - technically your shields are not up inasmuch as you can be seen. That is, but all that somebody knows is that, at that IP address, they don't know it's you, but at that IP address which your ISP has probably assigned you and it drifts around every few months, there is something that is there, and when you ping it, it says yeah, and when you give it a UPnP probe, it says yeah. But it's rejecting the UPnP probe, so it says I'm here, but I'm not listening to these.

Now, notice also that, if I had probed 1899 rather than 1900, 1899 would also reject. So the point is that it's not just 1900 that is rejecting an incoming UDP. It's all of your ports are rejecting an incoming packet, and I'm just happening to check for port 1900. So if you wanted to be stealthful, on an AirPort, you can do what we talked about a couple questions ago, which is set up a DMZ to a nonexistent IP, to an IP in your network, so it's like 192.168, but normally your router will assign maybe like from 1 to 100. So set the DMZ to 111, outside of the auto assignment range, so no machine will ever be assigned that IP, yet incoming traffic will go there. And when you do that, you will be fully stealth.

So you can be stealth. But you don't have to. Remember that all this tells anybody is some unknown piece of equipment is at that IP address. The theoretical vulnerability is they could, like, then try to dig deeper and do something to you. But at this point it seems unlikely.

Leo: Could you DMZ it?

Steve: Yes, that does work.

Leo: That's clever. I don't know if I've ever seen DMZ in an Apple router. I'm sure...

Steve: Oh, they must have it because that's the only...

Leo: They have to.

Steve: It's the only way to run a server behind your network. In fact, I should mention that's why the DMZ was created, not because of the Romulans, but because, if you wanted to allow - if you wanted to serve something, like your own little personal website at your IP, you inherently need unsolicited traffic, people you haven't invited by sending traffic to them, which normally means solicited traffic is coming back in response. But clients of servers are unsolicited by definition. So their traffic coming in needs to get to your server. So this DMZ is set up so that unsolicited, unexpected traffic coming in from anywhere is routed to that machine's IP, or that it can decide what it wants to do with it. So I'll bet that the AirPort has to.

Leo: Oh, yeah, no.

Steve: It's crazy.

Leo: In fact, MacBreak's telling me of course it does.

Steve: Ah, yeah.

Leo: Last question, our bonus round. Here we go, Q11. Jonathan in Toronto brings us the Software Utility Suggestion of the Week: Hi. Just came across a new little utility I'd like to share. You mentioned a few neat utilities a long time ago, KatMouse among them. I love KatMouse. It's been great, aside from the middle-mouse function that I would always have to disable on new installs. I also use AutoHotkey - I like that one, too - to change my Caps Lock into Control, and to reverse mouse scrolling since my drivers don't have that feature. The scroll reversal would stop working at random times, though. Well, since I now have a keyboard with hardware control of control - with hardware control of control?

Steve: Of control...

Leo: Oh, control key.

Steve: Right.

Leo: Via DIP switch setting, I no longer need AutoHotkey or KatMouse. Welcome WizMouse. WizMouse by Antibody Software has the familiar scroll-anything like KatMouse, but also has the option to reverse scrolling. Super simple, just a few options, just what I need. Hope you find it useful.

Steve: So I wanted to thank Jonathan for that. I have not looked at WizMouse, although I just did install KatMouse, both on the server at GRC and on my new Win7 box that I

was building to set up a clone of IIS so that I can do development and work here. And listeners may remember that what KatMouse does is it causes your mouse wheel to scroll whatever the pointer is over, even if it doesn't have focus. So focus is when you click on the window, and it darkens it and jumps to the front and so forth, which is normally what the mouse wheel will scroll. Even if you wander the mouse off of that to something else, and you scroll the wheel, that other thing still scrolls. KatMouse causes the mouse wheel to track with whatever the mouse cursor is floating over, which is really cool. I'm still using it.

I'll check out WizMouse. I don't have a problem with the scroll wheel being backwards because I like it the way my grandfather used it instead of the way Apple has wrecked everything by thinking that I'm using a touchpad where I want to, like, push it on my Mac the same way I do on my tablet, which is not the case. I find that the wheel is backwards on my Mac. So, and whenever I do a major update, I have to go in and reverse the direction on the Mac. But they do allow that to be done, so that's cool. Anyway, WizMouse. So check it out.

Leo: Good, good, good. And that concludes all 11 questions on this episode. Nicely done. Right exactly on time and everything. Steve Gibson's at GRC.com. He didn't mention SpinRite. That's what I was waiting for.

Steve: Yeah, I knew we were going to have a super - oh, that's what you're waiting for. I knew we were going to have a super long one. We had a lot of stuff at the top of the show.

Leo: All right.

Steve: And I thought, eh, everybody knows about SpinRite.

Leo: I'll mention SpinRite. Get it. World's best hard drive maintenance and recovery utility. You've got to have SpinRite if you've got a hard drive. And you can get it from Steve, GRC.com. Make that yabba-dabba-do sound off in his pocket when you buy a copy of SpinRite, really great thing to have. But while you're there he also does so many other free things. In fact, you might want to test that UPnP flaw using ShieldsUP!. So many other things you can do there. Just go to GRC.com. He also has 16Kb versions of this show for the bandwidth-impaired, and even transcripts, if you like to read along while Steve talks, thanks to Elaine.

Steve: Elaine sent me a note when she transcribed last week's podcast. She said, "It cracks me up so much hearing you describe where I apparently live." Because we were talking about that she has bats and...

Leo: Mountains and - there it is, right there. It's in the middle of nowhere is where she lives.

Steve: She did actually send me a photo, which I didn't focus on, but I showed it to Jenny. And it showed snow on a cactus.

Leo: Okay, there you go. Says it all.

Steve: That does.

Leo: [Laughing] Anyway, GRC.com. Also if you want to ask a question for our next feedback episode, two episodes hence, you can do that there, GRC.com/feedback. I think that's everything. We do the show Wednesdays, 11:00 a.m. Pacific, 2:00 p.m. Eastern time, 1900 UTC on TWiT.tv. Do watch live because then I can look at the chatroom and add addenda from the chatroom.

Steve: And we do have fun for the first 30 minutes, Leo. We ought to tell our non-live listeners that we normally get finally underway at about 11:30, but we have fun for the first half hour.

Leo: There's stuff going on. It's not like just dead air. You get to watch the studio set up.

Steve: No. It's a little bit - I'm a little conscious of the fact that I can say things that aren't going to be recorded for posterity. So, yeah.

Leo: Oh, well. Maybe not recorded. We can't promise that. People do record things, you know. So that's why you should tune in. But if you can't make it, we make on-demand audio and video versions available, as well. Just like Steve. Ours are a little higher quality at TWiT.tv/sn.

Steve: And there is a YouTube channel now, YouTube.com/securitynow...

Leo: Yeah, isn't that great.

Steve: ...will bring you the videos.

Leo: Yeah. If you want to watch the video on YouTube, that's good, too. Well, thank you, Steve.

Steve: Okay, my friend, we'll do a full episode of some cool topic, whatever is next on my list or happens in the meantime, next week. We'll see you then.

Leo: See you then.