Can "Anonymous" Take Down the Internet?

Description: This week, after catching up with the week's security and privacy news, Steve and Leo examine the feasibility of the hacker group "Anonymous" successfully taking the Internet offline after a disavowed Internet posting has claimed they intend on March 31st.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-341.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-341-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here, and we are going to get in a fight over Google, Apple, and the iOS cookie incident. That and a look at why Anonymous may or may not be able to take down the Internet next month. It's all coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 341, recorded February 22, 2012: The Anonymous Threat.

Time to protect yourself online with - boy, there couldn't be a better time for this show - Security Now!. Steve Gibson is here, the Explainer in Chief. And, well, first of all, welcome, Steve. Good to talk to you.

Steve Gibson: Thank you, Leo. Our recorded audience don't know that you and I have spent the last 30 minutes talking about coffee, before we pressed the Record button.

Leo: And we're feeling so good now.

Steve: And we're a little self-conscious about whether they would not have enjoyed our exploration as much as we did. But, oh, well.

Leo: Let us know. Let us know. Tell us. More coffee or less? I tell you, the chatroom is going crazy with recommendations and suggestions because there's something about geeks and caffeine. It merges together. And it all started because I saw Steve's ginormous cup. Look at the size of that thing. And I thought, somebody's got some coffee today.

Steve: And I love it, Leo.

Leo: There's something about coffee. And I believe we are now of the opinion coffee is good for you, not bad for you.

Steve: Oh, it is good. Not only the caffeine, but it turns out now they're finding that, because coffee bean has a plant origin, the plant-based chemicals are believed to have additional health benefits aside from just the pure stim effect, the so-called polyphenols that the coffee contains.

Leo: Now, this was a week, boy. The last three weeks, I would say, have been not security - sometimes we get weeks where there's just lots of hacking and stuff. And then this week was all about privacy.

Steve: Oh, my god, have you seen all this Google stuff happening?

Leo: Yes, yes.

Steve: We've got it all covered here this week.

Leo: We'll talk about that. And then there's a huge story, the folks at Anonymous - whoever that it is. It could be a guy named Jason who lives in Surrey, we don't know. But whoever that is has said they're going to take the Internet down.

Steve: A posting appeared on Pastebin, which is a way of anonymously putting things up and then sharing a link, which I will share with our listeners, stating the intent by the infamous Anonymous group, who everyone takes very seriously now. In fact, I've got a really interesting story from the director of the National Security Administration, the NSA, worrying about Anonymous's growing capabilities. So this is not something anyone ignores when they say this. They knocked off Visa, MasterCard, and PayPal during the whole WikiLeaks deal. When those payment processors announced that they would stop processing payments for WikiLeaks, Anonymous just took them down. They were blasted off the Internet for a substantial length of time. And RIAA has had continuing problems and so forth.

So they're now saying they're going to take the entire Internet down on March 31st. So I thought, well, let's not wait to talk about this because we have more than a month now. And so then there has been some follow-up with people claiming to be Anonymous, saying, no, that's not us. And of course the problem is, if you're truly Anonymous, you have no way of, I mean, I guess there ought to be, like - they ought to use PKI. They ought to have a way of signing their thing so that it could be proven that it's them. I don't know.

But anyway, so even if this was a bogus posting, trolling, as they have said, not from them, it doesn't make the question any less significant and interesting, and that is, is it possible to take down the Internet? And they're talking about doing a denial of service on the Internet's root servers, the root DNS servers, which is arguably the single focus point

of the Internet. I mean, we all need DNS in order to, as we know, resolve domain names into IP addresses. So I thought this was a great occasion for us to examine the technical feasibility of launching a denial of service attack against the Internet's root DNS servers. That's what we're going to talk about today, can Anonymous take down the Internet, if the wanted to.

Leo: I guess there's no point in saying who the heck are they, or is they, or what is Anonymous.

Steve: I think they call themselves a "hive." One of the postings I saw in Twitter with them saying, "That's not us," was them saying, "The hive is not going to do this. This is a bogus posting." It's like, oh, okay. Well, be a hive.

Leo: But that's the problem. How do we know? Who is "they"? I mean, if you're anonymous, who knows? All right. Well, let's get to that right now. I have a Ford commercial, which I will do sometime in a few minutes. I do want to kind of - you have some other things to talk about.

Steve: Oh, yeah. In fact, two important things relating to last week's discussion of the discovery that there was so much collision among the prime numbers that were being generated, which are the basis for public key technology on the web. Well, we'll remember that an academic analysis was done by a group of researchers who looked at the public keys and saw a surprising level of collision among them. Well, first of all, I don't know what I was thinking, I just misspoke, which is just a classic brain fart, is what you'd have to call it, where I talked about how the certificate authorities had to do a better job of generating prime numbers. Which is not how it works at all.

How many times have we talked about the fact that what they do is sign our public key? We generate the public and private key. And the essence of the elegance of the system is that the private key, which is to say the prime numbers which we choose, never leaves our control, never needs to. We generate the public key from the two prime numbers. We send that off along with verifications of our identity to the certificate authority that then signs the public key, thus blessing it and proving that we're the people who have the public key.

So this actually makes the problem a little more worrisome because it means there isn't a single point of, like, heightened responsibility where we can say, oh, well - as I, in a bonehead fashion, did last week - like all we need is for the certificate authorities to do a better job generating the certificates. They don't. All they do is sign the ones we provide them. Which means that what this tells us is that, I mean, the individual servers or machines that people who generate the so-called CSR, the Certificate Signing Request, which is what's sent to the certificate authorities, all of those machines are where we need them to be doing a better job of choosing random prime numbers, which are then aggregated, multiplied together in order to generate the matching public key. So that's like, well, okay, that's more of a problem. So that's one.

The second thing that I missed, and actually a buddy of mine, an online friend, Paul Byford, who goes by the handle "Sparky" in our newsgroups and has become a sort of long-term contributor, he said, "Steve, you forgot about the Euclidean algorithm aspect of this." And it's like, oh, of course I did. And that needs to be discussed, mostly because it's just so cool.

So Euclid figured out, like in 300 B.C., that there was a cool way of finding the greatest common divisor, the GCD, also the greatest common factor, the GCF, same thing basically, of any two numbers. And I really recommend that our listeners go take a look at Euclid's algorithm on Wikipedia because Wikipedia has a very nice treatment of it. This is not rocket science. You don't have to have a degree in math. What's so cool about it is how simple it is. It's sort of a construction-based algorithm. Remember, this is 300 B.C., so he didn't have access to advanced math...

Leo: He had rocks.

Steve: ...and all that. And remember, like, construction things, like when we were learning in geometry how to form the chord of a circle, or how to inscribe an isosceles triangle in a circle, where you basically just sort of do things with a compass. This is sort of like that. The idea is you have two numbers - and on Wikipedia they have a beautiful little diagram where they sort of explain graphically how this algorithm works, where you take the shorter, if you think of the numbers as lengths, you take the shorter of the two and subtract that from the longer until the result of the subtractions is smaller so that you can't subtract anymore. Then that gives you a certain length. Then you apply that against the other ones, similarly, and you go back and forth until you can't go anymore. And then that is the way you're able to find the greatest common divisor. It's just this incredibly simple, very clever system.

But think what that means in terms of public keys, which are the product of two primes. It means you don't have to find two public keys that are the same. All you need to see is whether two different public keys have a common greatest common divisor, which is prime. So, and since the public key is a product of two primes, you basically take all the public keys, and you apply Euclid's algorithm against every pair, even if they're different. So it's not that they have to be the same. It's that, if two different public keys share one prime number, this algorithm instantly, I mean, easily finds the common prime because that will be the greatest common factor that these two public keys share. And once you've got that prime, of course, you divide that back - you divide the public key by one factor, and that gives you the other. So you can crack the private keys of both websites if they happened to share a common prime.

And thus the title of the paper. The title of this academic paper was "[Ron was wrong, Whit is right]." And so what their point was, okay, Whit Diffie was the guy who came up with a public key system that only used one secret. Ron Rivest, who's the "R" of RSA, he uses the two primes. The point being you're in much worse shape if you have two secrets, meaning these two prime numbers, than if you only need one secret. And so what they were saying was that the fact that it is so trivial to take two products of primes, which is to say the two public keys, and almost instantly find their greatest common factor, dramatically weakens the fundamental security of our public key infrastructure in the situation of there ever being a chance of sites using the same prime.

And what's been found was sites are stumbling on, by pure chance because the random number generators are not very good, they're stumbling on the same primes, multiplying them together, and then, turns out if you analyze a large database of public keys, you find many more collisions than was ever expected before. So that is really cool. Not good news, but just such an elegant hack, essentially, on something that we assumed was going to be much secure than has turned out to be. Okay. So Google. Whoo, boy.

Leo: Whoo, boy.

Steve: Oh, boy. One thing I needed to say was that Google is working on a password generator for Chrome.

Leo: I saw that. I would presume it's something like SuperGenPass.

Steve: It's like LastPass

Leo: Oh, you mean they're going to be a password storage thing, too.

Steve: Yes, yes. They will generate and store and cloud sync. So it's the whole LastPass enchilada, essentially, that will be built into Chrome because Google recognizes that, I mean, in their own blog posting about this, they say, okay, so passwords are not what we wish we had. We need something better. But they're all that we have at the moment. So let's make them as good as we can in the short term. So essentially they recognize that LastPass is a great solution, and they're going to build something like it natively into Chrome. It's not happened yet, but they're exploring it. And it looks like they're going to do it.

So, yeah, I'm with you, Leo, for dyed-in-the-wool Chrome users, that's only good news because I guess what I like about it is that it will expose this technology to a much larger audience. I mean, if everyone in the world were using LastPass, problem solved, largely. But everyone in the world's not, and everyone in the world's not going to. If it were built into Firefox, then Firefox users would all have it. And when it gets built into Chrome, all Chrome users will have it. And it'll just end up increasing the overall security on behalf of the users of Chrome because it'll just be there. And we know, from all of our experience, that what's there matters. Things that are there get used. Things that you have to go out and find and install and then configure, eh, not so much, unfortunately.

Leo: I just thought, when I read it, and I didn't read it very closely, that it was going to be something like SuperGenPass, where it would generate secure passwords for you. But I guess it would have to remember them, or what good would it be?

Steve: Yeah, exactly. And cloud syncing is the other thing that LastPass gives us that'll be there, too. And so it'll be across all Chrome browsers on all platforms.

Leo: Right. I mean, I use Chrome everywhere, and I use their syncing capabilities and their password syncing already. But I also use other browsers and other platforms and mobile platforms. So I think LastPass is not going to be out of business.

Steve: I agree. That's what's really nice about LastPass is I've got it over on Chrome, and suddenly Chrome knows all of my LastPass passwords, too.

Leo: Exactly.

Steve: Speaking of which, somebody tweeted - I'm jumping out of order here, but we're on the topic, so I will. David Ward, who tweets from Sydney, Australia said, "Re giving loved ones access to your passwords when you pass on, what about a LastPass OTP left in a safe? I think that works." And I'd forgotten that completely. That's a perfect solution.

If you Google "LastPass OTP" or "one-time passwords," remember that another feature of LastPass is that you can generate one-time passwords which they save encrypted in their servers. And what they do that for is to solve the problem of needing to access LastPass in a public setting, like on a library computer, for example, although I don't know that there is any safe way to do that. But still, if there were a key logger, for example, which could log your keystrokes, you would be ill advised to log into LastPass using your normal, everyday LastPass master password because the keystroke logger could glom onto that, and your security would be breached.

So you can pre-generate and carry in your wallet a series of one-time passwords for LastPass to solve that problem. You log in once using it. It will allow you to then access your LastPass account, but you will never be able - in the act of using it, you tear that one up. You can never use that one again. So that's a perfect solution. You generate one or more and give it to your attorney, stick it in your "When I die" bank vault or whatever, and that's a way for your loved ones to get access to your accounts. And then that solves the problem of, and this was the point that was brought up last week, was the regular user wants to be able to change their LastPass master password periodically for security, which is perfect. You can still do that any time you want to. That does not obsolete your one-time passwords, which will live on until they're used. So David, thank you for bringing that up. That was a great tweet and a great solution to the problem. Okay. Back to Google.

Leo: That's the good stuff. Now the bad stuff.

Steve: So, okay. There were two main issues that were brought up and got wide coverage. Jonathan Mayer, who we've spoken of before, is a researcher at Stanford University. And he discovered something rather disturbing, which was that Google and three other advertisers were deliberately bypassing Safari's anti-tracking. Safari is, and I love Apple for this, and I love Safari for this, the only browser that has third-party tracking cookies disabled by default. It's just amazing to me.

One of the last things I still need to get to, and it's on my very short list of things I need to wrap up in order to get on to my next big project, is my cookie system. If you go to GRC.com/cookies/cookies.htm, Leo, you will see a still-unpublished set of pages with technology that's been actually running now for several years at GRC. Last week GRC had 73,356 unique visitors. Of those 73,356 unique visitors, 84.59 percent of them had third-party cookies enabled. So almost 85 percent. Think about that. I mean, the vast majority of people surfing the 'Net still have third-party cookies enabled. Why? Because that's the default setting in all browsers except Safari. And if you scroll down that page, I'm tracking all of the cookie usage by browser.

Leo: Are you using cookies to do this? No, I guess you couldn't.

Steve: Oh, yeah. Oh, yeah, yeah. I've got some amazing technology. If you look at, in the block of links at the bottom, look at the stats. I think it's the top link on the third column because I clicked it just before this as I was generating the URL. This is showing it dynamically tracks the statistics of cookie usage by browser version and manufacturer. And there's a bar chart, if you scroll down a little ways, showing how different Safari cookie usage is than every other browser. Why? Because Safari has them off by default. Only 26.65 percent of Safari visitors have third-party cookies enabled because it's off by default. So that's, I mean, that's a radical difference, demonstrating, once again, the importance of default settings. Most people are going to just be using default settings all the time.

Leo: You and I have kind of a disagreement on this issue.

Steve: Ah, okay.

Leo: Well, I don't think - so perhaps you might explain why third-party cookies are bad.

Steve: Well, oh, no. And we don't disagree, actually, as much. One of the things that we will get to, and I'm glad you brought it up, is okay, so what? How much does it matter? Is it a big deal, and so forth.

Leo: In fact, a lot of sites won't work. We don't use third-party cookies particularly on our site. But occasionally we do, if you have an ad banner. What a third-party cookie is on our site, you come to TWiT.tv, and there's an ad banner there. That ad banner can set a cookie on your page and use that cookie to determine that you visited the page and so forth. And that's how we get paid. And in fact that's how - third-party cookies are in fact how the Internet monetizes. So I think you could make a case, and John Battelle made this case, that Apple in fact isn't doing this for privacy purposes, but to, as Apple is wont, keep others from profiting when it's Apple's job to make all the money on Mac users and iOS users.

Steve: That's interesting.

Leo: Anyway, I'll just bring that up.

Steve: Okay, so...

Leo: Go ahead, what's wrong with third-party cookies?

Steve: Well, okay. So, arguably, I think the user should be informed.

Leo: I don't disagree with you on that. I completely agree with you on that.

Steve: And so I would have...

Leo: But setting a default, in a way, is not a form of information, it's merely blocking it on their behalf, determining what they want.

Steve: Agreed. And so, for example, if third-party cookies are important to a site for its own monetary support, I would have no problem with the site saying, "Hi there, welcome. We're happy to have you, but you've got third-party cookies disabled, and we need those in order to get paid. So if you want to use our site's free services, click here to turn on third-party cookies so that our advertisers can know where you came from and can pay us for the fact that you're browsing around our pages."

And, I mean, if we had a system like that, and it's entirely doable, but it just hasn't happened yet, then I think the problem is solved. It would be built into browsers. Yes, it would get in people's face, but be much gentler than NoScript, which is the page doesn't work at all, and people are wondering why and so forth. But it just says, for those sites that want to, to ask their visitors to turn on third-party cookies.

Now, users could then decide how they feel about that. Savvy people might say, well, you know, I don't really need to be here, and I'd rather not be tracked because the concern that people have had is that third-party cookies, the whole point of them is that it's not the site that you're visiting, it's the site whose resources the site is presenting, like ads, that then allow profiles to be made of people over time. I know from my own experience that there are people who feel very strongly, I mean, they delete their cookies, they flush their cookies, they go through all this work. They just don't like the idea that they're going to be tracked. It's like, okay.

So you and I are on the same page that users ought to be informed. And I have no problem with a site saying, "Hi. You don't have third-party cookies enabled. We need them in order to be paid by our advertisers, so please turn them on for this site." And that technology exists in all the browsers. You can do per-site enabling of third-party cookies, which solves the problem. Right now, it's just not easy, and it ought to be made easy.

Leo: Yeah, I mean, I go back and forth on this. I think that most, the vast majority of people have no idea what a cookie is. And unfortunately, I think that there's a lot of paranoia building about cookies. I don't think cookies are harmful in the least. In fact, I think they're how the Internet works. And I think that blocking cookies, first or third-party cookies, breaks the Internet. And people prey on the, I think, incorrect paranoia people have about cookies. People delete cookies like crazy, as if that's some magic process. It is not. And as we have talked about before, it also does not in any way prevent companies from tracking you.

Steve: I was going to say, exactly, and there are many other ways of locking onto someone. Still, cookies are the main way. They're the easiest way. They're the most straightforward way. And you could argue, I mean, they were built into browsers in order to create this sort of state, in order to create a stateful connection to visitors. So, yeah. I guess, what, long-term, in the same way that scripting is something that you can increasingly not live without, cookies are probably something that you will increasingly need to allow.

Leo: Turn off cookies and see how the Internet works. It doesn't.

Steve: Yeah.

Leo: The good news is you can turn off third-party cookies. Apple calls them "Allow cookies only from sites I am visiting."

Steve: Right.

Leo: And you can in fact disable that without any apparent hardship. Although there are some sites, somebody in the chatroom said the BBC site will not show videos without third-party cookies. There are some sites that does break. But because this is so prevalent on Apple hardware, most sites have found ways around it. And that's what Google was doing, and these other ad agencies.

Steve: Precisely.

Leo: The question is, why does Apple do that? Are they on the side of the angels? And it's merely my opinion that Apple, they do this all the time, they don't want anybody to have information about their users except Apple. So that's the real issue. Apple knows exactly where you're going and serves you ads in response.

Steve: There could be a real reason. They actually do say it's for user privacy.

Leo: Well, of course they do. That's what I would say, too, were I Apple.

Steve: Okay.

Leo: That doesn't mean that's why. It just means what they say.

Steve: The way their cookie system works softens this a little bit. And this is what I learned from this cookie forensics experiment that I did years ago because this thing characterizes the exact cookie handling of all browsers. And it turns out no browser is free of actual bugs in their cookie handling, some worse than others. In Safari, if you disable third-party cookies, it doesn't simply kill them or stop them. If you had any third-party cookies, they continue to be sent to any third-party site that you visit. So it won't accept new ones, but it still transmits any that you have that match that domain. If you are at a third-party site, and you have a cookie for that site, that one site has the ability to write cookies to you. So a third-party site that your browser doesn't currently have a cookie for cannot write new cookies. But a third-party site that your browser does have a cookie for, first will receive it, and can modify it, is able to write cookies. So in order to, if you were concerned, you'd have to disable cookies in Safari, delete them all, then restart Safari because there's also a whole restarting thing that comes into this, whether they're,

like, many browsers still don't obey your requests until you shut them down and start them up again.

But the third interesting aspect is, in Safari by design - now, this was removed from WebKit seven months ago, but it hasn't yet migrated into Safari. And that is, if you submit a form to a third-party site, then that third party is allowed, even if third-party cookies are disabled, the response to a form submission to a third-party site is allowed to submit cookies. And that's what Google figured out.

Leo: So you're saying that WebKit disabled this months ago.

Steve: Yes.

Leo: And of course Safari is based on WebKit. But Apple did not.

Steve: Apple hasn't. I don't know if Safari is tracking the current WebKit. But Google took it out of WebKit seven months ago.

Leo: Google took it out.

Steve: Yes.

Leo: So WebKit is an open source project that many companies contribute to, including Google and Apple.

Steve: Right.

Leo: And their browsers, both Chrome and Safari, are based on.

Steve: Exactly.

Leo: But it was Google that removed that from the WebKit code?

Steve: Seven months ago.

Leo: In Chrome, or in the WebKit open source project?

Steve: In the WebKit open source project.

Leo: So if you use the WebKit browser - see, the problem is on iOS you can't use the WebKit browser. But if you're using the WebKit browser on the desktop, this trick wouldn't work.

Steve: Correct.

Leo: Interesting.

Steve: But it works on Safari because Apple hasn't been pulling those changes from WebKit.

Leo: Interesting.

Steve: So it's still there in Safari. So what Google figured out, and three other advertisers, as well, and this is what Jonathan Mayer at Stanford University caught them doing, is they needed to bust Safari's blocking - Google "they." Google needed to bust Safari's blocking of third-party cookies for their +1 feature to work on ads because ads are from third parties, and they needed to essentially enable third-party cookies. But they couldn't in Safari by default. And it was a problem because Safari has this disabled by default. So as we've seen, more than three quarters of Safari users have third-party cookies disabled.

Google didn't like that. So what Google figured out was - get this, Leo. A web form in an iframe in an ad can be submitted by JavaScript. And that allows them, through this really convoluted mechanism, to get a third-party cookie set. And what Jonathan discovered was that Google's code looks at the user agent. And he tested 400 different user agents and has a spreadsheet available in his original posting showing this, that only when the user agent is Safari, any variant of Safari on any of the Safari platforms, then different code is issued which uses this iframe form posting trick, which is then submitted so the user has no interaction needed because JavaScript triggers the submit function of the form in order to get this +1 functionality to work.

Leo: That's, by the way, very common practice in JavaScript. You test user agent because JavaScript itself has so many bugs. So this is not - it sounds like, again, I think it's important not to overblow this. It sounds like, oh, my god, they were searching for user agents. It's throughout all JavaScript code that you do this because you always have different code for different user agents. That's unfortunately a necessity because JavaScript is so buggy. And what did Google do with this end-around? Were they tracking people as they surfed the 'Net? No. They turned on a +1 button for people who were logged into Google.

Steve: This was allowing, exactly, this was allowing third-party cookies to be set.

Leo: Right. And as far as we know, it's possible they were doing other stuff, including tracking people's iPhones, I guess. But it was, in my opinion, I think

Google's justified. But again, there's two sides of this story. I'm just trying to give you both sides. What Google's doing is, for people who are logged into Google, is to turn on the +1 button so that you may +1 ads as well as +1 pages. That does not work, just as the Facebook Like button does not work, if third-party cookies are disabled.

Steve: Right.

Leo: And one could say, in fact, as many do, the Like button and the +1 button are invasive on all web pages because, as soon as you go to a page with a Like button or a +1 button, unless you've disabled third-party cookies, Facebook or Google, respectively, know you're there. In fact, we've even talked about it. Even if you're logged out on Facebook, apparently, the Like button still sends a signal.

Steve: Right.

Leo: So the real question, I guess, is how nefarious - it does look pretty sleazy to end-around, to look at a workaround for something like this. And I think Google probably should stop doing that.

Steve: Yes.

Leo: On the other hand...

Steve: And they have said they're going to. They have said we're sorry, we'll figure out some other way around this.

Leo: Right. But on the other hand, companies like Facebook and Google provide some significant services for free. And the way they monetize is this way. And so it's kind of how - it's kind of how that works.

Steve: Yup, the reason it's free. Now, to their discredit, Microsoft jumped on this. And this is Part 2 of this brouhaha. Microsoft said, well, Google is abusing our P3P technology, which is Microsoft's Platform for Privacy Preferences. Quoting from The Verge, which is an online news source, they said, "Just a few days after the Wall Street Journal reported that Google, Facebook, and others have been using a workaround to bypass the cookie restrictions in Apple's Safari and Mobile Safari web browsers, Microsoft now claims that Google has taken similar measures to bypass privacy settings in Internet Explorer. Microsoft says that Google is improperly representing its cookies by using a non-standard P3P cookie policy statement. It claims that 'Google's P3P policy is actually a statement that it is not a P3P policy,' which allows Google's cookies to pass through without being blocked."

Google's response to this is also reported by The Verge: "Earlier today, Microsoft accused Google of manipulating Internet Explorer's default privacy restrictions in order to 'bypass

user preferences about cookies.' Google has just responded with a lengthy rebuttal, arguing that Microsoft's P3P cookie technology is 'widely non-operational,'" - love that phrase - "and that the issue has been around since 2002. The response also points to other offenders, citing a 2010 Carnegie Mellon research paper that says over 11,000 websites don't use valid P3P policies." And then Google talks about Facebook and Amazon, saying that they also are doing the same thing.

Now, in this, I completely agree with Google. What happened was Microsoft, in an earlier version of IE, entertained disabling third-party cookies also by default. There was a version for a while in beta that had third-party cookies disabled by default. Microsoft generated so much flak from doing this by big business that just screamed. And this is years ago. This is 10 years ago or so, and I don't know, maybe it was an early version of IE6. But Microsoft got so much flak that they backed off from turning off third-party cookies by default and came up with this bogus approach, where a website, in the query headers going to a server, can assert what their cookie policy is in a machine-readable header of little three-character tokens.

And what really annoys me is that, if you, in IE today, all versions of IE, if you turn up your privacy to maximum, such that it says you are blocking third-party cookies, yet if a third-party website has a specially crafted P3P policy header, IE goes, oh, well, they say that they're going to do good things with your cookie, so we'll let third-party cookies work anyway. And so this is a complete override over IE's clearly stated policies that the user can control, even on a per-site basis. If a site says, no, no, we're nice people, then IE just says, oh, well, in that case, let me have your cookie. So that's what Google is saying. They're saying this is ridiculous. IE allows any site that wants to, to override the user's and the browser's preferences. So we're doing it, yes, and so is Facebook and Amazon and 11,000 other sites. So that's what that was. What do you think about that, Leo?

Leo: Interesting. Well, it's not a surprise. I mean, look, everybody wants to make as much hay as they can out of this because - the real issue, the very significant issue, unfortunately, is this is highly technical stuff, and there's a lot of issues involved that are very complex. And even the tech press is - it makes a great headline.

Steve: They're never going to get it right.

Leo: It makes a great headline. And so people are - I get a lot of accusations whenever I bring this up, oh, you're biased in favor of Google. And I think anybody knows who listens to everything I do that I don't have a bias in favor of Google at all, or Apple. But I think what has to happen is that, as best we can, to this audience which is highly technical, is for the audience to understand not merely how has it happened, but there are whys.

And there's a deeper issue, which is how the Internet monetizes itself. And as someone who, I mean, I'm kind of above this because we monetize by ads, which we force you to watch, and I'm about to do one. But you have no choice. So we do have some banner ads, and we are going to start doing more banner ads on the website, and that's where all this starts to happen. That's when all this third-party cookie tracking and so forth starts to happen. It is the means that, frankly, powers the Internet.

And the good news, Google, I think, is extremely upfront in their privacy policy. And the good news is, if you do not log into a Google account, or if you do not create a Google account, or you delete your Google account, none of this will happen to you. You do have control of this. This only happens to people who are logged into Google at the time that they're using the browser. And why do you log into Google? Because you want to use their free services, chiefly Gmail, or perhaps Google Plus or other services. You can use Google search completely anonymously. It's a completely free service. They do not monetize that directly by tracking you. They put ads in the search results, but they do not track the results unless you log into a Google account. And I think they've been remarkably upfront about this in their privacy policy, which they urge people to read over and over again.

Steve: I guess the part that I don't get, though, is why third-party cookies have anything to do with monetization.

Leo: Well, a perfect example is the +1 button.

Steve: Okay. But that's new. So, for example, if you host ads on TWiT, then...

Leo: There's a third-party ad on the banner site; right? Those ads are not delivered from TWiT. They're delivered from an ad server, which is a third party. How do they know that ad was viewed? Oh, they get hit by the IP address, that's right.

Steve: No, the referrer header says we're serving this ad to TWiT.tv. And so that's where you get your impression count. The fact that a cookie wasn't sent with it only really means that this user doesn't have their cookie.

Leo: That's right. That's a good point.

Steve: So they would like to know who you were because then they could aggregate all the other places you have been and everything they know about you, and maybe serve a higher value ad to you. So it's really a - its argument is it's a user benefit because, if you are profiled by the third party, then they know you're 75, based on the fact of the medication that you've been looking at. And they're not going to give you a diaper ad, which isn't relevant to you. And so the idea is...

Leo: Well, there's value to both sides. There's value, of course, to the ad server because they can sell that ad for more money because it's a targeted demographic. So I think the real issue - and I don't think most people would have a problem, or having understood that, I don't think that is - because they're not tracking you personally. It's not like they say, how can we get an ad to Leo Laporte? I don't believe that ever happens.

Steve: Well, yeah, the argument is that they really do know who we are, that there's enough privacy...

Leo: They don't care. But that's not - what they really want to know is what's my age, income; they want to know demographic information. That's what's of value to them, not my name, age, and social - not my Social Security number, you know what I'm saying? They don't care about me as an individual. It's an aggregate. At least my understanding of how ads works. Now, the real question for people who are worried about this is exactly that. Are they trying to find out something about me personally? And I don't believe that's the case.

Steve: Well, people who have worked in the third-party advertising world have said you wouldn't believe how much information they have about people.

Leo: Of course. In fact, there's a great article in the Sunday Times which I recommend everybody read about how Target knew that a teenager was pregnant before her parents did. And it's not, by the way, necessarily online behavior. They're aggregating tons of data about you.

Steve: Pulling it all together.

Leo: But again - and I think part of this is a disconnect in how we interpret the word "privacy" in the Internet age. It isn't really that they cared about that girl and they somehow wanted to know is she pregnant, that particular girl. What they want to do is identify second trimester customers so they can target particular kinds of advertising, its higher value ads, for a lot of reasons. Read the article, it's fascinating. And it doesn't have anything to do really with online privacy, it's just in general. Is that an invasion of privacy? It's not like there's a guy at BBD&O who says, hey, Steve Gibson loves a certain kind of coffee. Let's send Steve Gibson - he doesn't care who Steve Gibson is. That's too small. That's too granular. That's not how you make money. Maybe someday.

Steve: Databases are huge.

Leo: Right. Maybe some day you'll make money selling individually to Steve. But it's much better if I know the three million people who have a burr grinder. That's what I want to know, not the one person. So I guess my question, again, is do you feel that people are trying to figure out what Steve Gibson is up to, like they want to know you?

Steve: What I know from lots of discussions in privacy newsgroups is just - it's a creep factor. When someone hears, for example, your story, that Target knew a girl was pregnant before her parents, some percentage of people are going to get creeped out by that. They're just not - they just don't like the idea.

Leo: And I understand that's your right, to be creeped out by and to protect yourself. Completely agree with that. And so my point is not that it's - you have to make the decision for yourself. But I want people to make the decision based on

absolute fact as opposed to the kneejerk sensationalistic headlines that is what is being distributed around. And for me, I don't think - they don't care about me. It's not like they're peering in my window. Yes, they know a lot of information. There's a lot of information about me in a database, with the idea toward sending me targeted advertising. I don't find that particularly intrusive, but maybe some do.

Steve: And just so we're clear, my fascination is just the technology. I get off on how all this stuff works. I love understanding it.

Leo: I agree.

Steve: And our listeners are saying, hey, how does this work? I want to know how it works so that I know what it means.

Leo: And it's our job, it really is our job just to get as much of the information, as technically accurately as possible, out to you. And then it's ultimately up to each individual what they want to do about it, and whether they're up in arms at Google or not, or all of that. So I just want - I only bring it up because I think that there is another point of view that needs to be expressed about what this stuff is so that you understand fully what they're doing with it and why they want it. And we get a lot of free stuff on the Internet, and that's why.

Steve: But it's not. I mean, tracking isn't really necessary.

Leo: Well, but it's necessary for a certain level of income.

Steve: Yes.

Leo: Now, I don't know how much it costs Facebook to give me a Facebook page. And so I don't understand the economics of that. They made $1 billion last year. Is that enough to pay for 850 million or users or not? I don't know. They can make a lot more money with more targeted ads.

Steve: Yes. I was just going to say, it is one thing we do know for sure, and that is that ad targeting does dramatically increase the value per impression.

Leo: And they may need that. It may not be sufficient just to get the click. That's all we get is the CPM. But it may be that $10 or whatever it is we charge per thousand impressions, that money may be predicated on the fact that the server, ad server, knows more about the viewer than I know. I don't know anything. They know it all. So I just know that we have a certain cost of doing business. We need to be able to charge a certain amount for ads. And I don't think it's for any...

Steve: And if the ads you show are...

Leo: Targeted, they're more valuable.

Steve: ...of higher value to the advertiser, then that's good for you.

Leo: Right. So I don't think we can say that it's not necessary. We just don't know. And you can opt out. There are plenty of people who watch this show who use ad blockers. There are plenty of people who listen to this show who don't listen to the ads. You're getting something for free that you're not paying for with your attention. And so I think this is the problem, is that I don't think people think of their attention as a currency, but it is in fact a currency.

Steve: Oh, Leo, I'll confess, I can't watch live television. I can't. If I didn't have the ability to jump over commercials, I'd just go crazy.

Leo: Right, right. Moving along, Mr. G.

Steve: So I did notice that Gizmodo picked up on the spray-on antenna story.

Leo: Ah, I was wondering if we'd hear more about that.

Steve: Yeah, I haven't - there's no additional information so far. I've watched a lot of background Twitter going around. They have patents. Apparently this, was it Cham something, Cham, can't remember the name of the site, but I don't think it's Sham.

Leo: It is spelled C-h-a-m, so there is some...

Steve: It is. But apparently they are a government contractor, and they've got a bunch of patents on stuff. So I guess we'll just kind of keep our antennae up, so to speak, and tell you if we hear more from them.

Leo: Chatroom says it's "ChamTech" like "chameleon." That makes sense, like "chameleon," because they're painting this. It's bizarre, bizarre, yeah.

Steve: Cool. Let's see. I found an interesting site I wanted to share with our listeners, BuiltWith.com, and also Trends.BuiltWith.com. What this is, is it's a search engine that goes out and inventories all of the sites on the Internet and looks at the technology that the sites are using, what versions of PHP, are they using Flash, are they using Shockwave, are they using RealMedia and so forth. And their trends page is really interesting because you can see things like the ebb of the use of Flash over time, where sites are, sure enough, moving away from using and relying on Flash to an increasing degree. So anyway, I just thought it was kind of a cool site that came across my radar

that I wanted to share with our listeners, BuiltWith.com.

And in sad news, Chrome has lost its side tabs. I went looking for them the other day. You could go to about:settings or about:config or something and turn those on. Well, apparently it was only just an experiment. They said on their blog posting about this: "As an experiment, side tabs were not a success. A small number of people really passionately loved them" - yup, count me - "but they ended up not being compelling enough to make the cut. We torture ourselves over stuff like this. It comes down to painful decisions about keeping Chrome lightweight. We know that a feature like this is really important to some number of users and Chrome developers!" - exclamation point - "but at the same time we have to continually cut and trim, knowing that those cuts will annoy people, so that Chrome doesn't turn into bloatware and satisfy no one."

Leo: Now, if Google had said "will annoy Steve Gibson," would you have felt your privacy was invaded? I think that's who they were thinking of.

Steve: They said, "We do hope to have a better solution to the 'I have too many tabs' problem someday soon, but side tabs wasn't it. I'm really sorry that we let the experiment linger too long. It meant that many of you became dependent on it, making the end of the experiment an even bigger pain than we wanted it to be."

Leo: They knew that there were people out there.

Steve: Oh, there was a lot of flak from that. And anyway, I also ran across something called Tabs Manager for Chrome that I wanted to give people a heads-up on. I don't know what happens when you have too many tabs on Chrome because I'm not yet a full-time Chrome user. I'm still over on Firefox with, like, 58 tabs open at the moment. Actually I do know that that is the number that I have open.

Leo: What, did you count them? Or is there a number there?

Steve: No, there's a tab session manager.

Leo: Oh, that's funny.

Steve: And so sometimes I'll just save all the ones I've got because I don't want a crash to cause me to lose them. I mean, I just use them as bookmarks. They're like things I want to get back to when I have a moment. When I surface from coding, I'll, like, read a few pages that I just didn't want to interrupt myself to read. So it's just like bookmarks on the Internet is in that fashion, is managed that way. Anyway, this Tabs Manager, two words, Tabs Manager for Chrome, it just puts a little button up on your button area of Chrome. And when you click it, it gives you a nice listing of all your tabs that are organized, and you can drag them around and see them easily. So it's sort of like having the tabs all there because it's easy then to click on one and access a tab.

So anyway, I'm glad - I want Chrome to solve the problem. If they come up with a better way than actually having tabs on the side, hey, that's fine, as long as there's some way

to deal with it because many people like myself organize our web browsing and our lives around mass quantity of tabs. So figure out how to do that, Chrome, or Google. That would be great.

Leo: My suspicion is Google is going to invent something besides tabs that they will say is a better way.

Steve: Yay.

Leo: That they're trying to solve this.

Steve: That would be - I'd be happy to have them do that. And speaking of happy, a subject, "SpinRite made my wife cry"...

Leo: No.

Steve: ...caught my attention. On February 8th, Andrew said, "Hi, Steve. I'm not a super tech, but I love the podcast for all the great information and news you and Leo provide. To get right to it, I have had a hard drive that apparently died on me about six years ago. Multiple attempts were made to recover the data as it had three years of family photos, including one of my sons' births and two years of his infancy. The only other option I felt I had was to send the drive off and pay a large sum to have the data recovered. I figured maybe in the future someday I could do this, as money was short currently.

"I've been a listener for a year now, and I figured, what the heck, let's try it. I purchased SpinRite and began the recovery process. The scan ran for approximately two weeks." And now I'll just remind people, that's like a worst-case scenario. SpinRite will work as long as it has to, to do the recovery. Normally it's two hours. But it can be two weeks if there's, like, lots of extensive damage. And he said, "I assumed 'It's probably not going to work.' When the scan finally finished, I connected it as a secondary drive to my PC. The drive appeared, and the data was now accessible!!!" Three exclamation points.

"I immediately copied over all the data. I put together a slide show of our precious photos and ran it. When my wife was passing by the computer and realized what had happened, the biggest smile and streaming tears came to her face. Thank you so much for saving us hundreds, if not thousands of dollars, and making a very memorable moment in our lives."

Leo: She cried in joy, not in sadness.

Steve: So Andrew - tears of happiness. So thanks for sharing that, Andrew.

Leo: We're going to get to the meat of the matter in just a second, the Anonymous threat. But I thought, given our discussion of privacy, I should mention that after

this show we're going to do This Week in Google with Jeff Jarvis and Gina Trapani. And our guest actually has written a book about this subject. It's called "The Consent of the Networked: The Worldwide Struggle for Internet Freedom." And one of the topics is, she says, "It's time to fight for our rights before they're sold, legislated, programmed, and engineered away." She's talking about privacy, among other things. And so this will be a good discussion. Rebecca MacKinnon will join us on This Week in Google in about 20 minutes, for those of you watching live. And for those of you listening, that would be a good one to download, if you want to hear more about this debate.

Steve: Very cool.

Leo: Yeah, fascinating subject.

Steve: Very cool.

Leo: Anonymous. They're after us.

Steve: Okay. So we don't take Anonymous lightly. They did actually take down the CIA.gov website earlier this month. On February 10th was the news about that. And when I saw the news, I immediately went to - I think I saw it in real time via Twitter, and I went there, and the site sure enough was down. And it was down for a while. The Associated Press on Friday the 17th reported that Anonymous had breached the United States Federal Trade Commission's Consumer Protection Business Center website, as well as a National Consumer Protection Week website. Both sites were temporarily replaced by a "violent German language video," focused on the Anti-Counterfeiting Trade Agreement, the ACTA that we've talked about. So, I mean, these guys are the real deal. And The Wall Street Journal on February 21st quoted the director of the NSA, the National Security Agency, warning about the growing strength of the group Anonymous. And in The Wall Street Journal article they wrote:

"The group has never listed a power blackout as a goal." That's what this NSA guy was worried about was that they were acquiring the ability to access the United States power grid and take parts of it offline. So The Wall Street Journal article continues: "The group has never listed a power blackout as a goal, but some federal officials believe Anonymous is headed in a more disruptive direction. An attack on a network would be consistent with recent public claims and threats by the group. Last week, for instance, Anonymous announced a plan to shut down the Internet on March 31, which it calls Operation Global Blackout." And of course, as I mentioned at the top of the show, they have taken Visa and MasterCard, PayPal and other payment providers offline and so forth.

So in this Pastebin posting, I created a short link, a bit.ly link using this Security Now! episode number. So it's bit.ly/SN341. That'll get you there, if you're curious. I won't read it in detail because I don't want to just take up the time for that. But essentially this appears to be a legitimate posting from Anonymous saying, "To protest SOPA, Wall Street, our irresponsible leaders, and the beloved bankers who are starving the world for their own selfish needs out of sheer sadistic fun, on March 31st, Anonymous will shut the Internet down."

Leo: Aw, you kids.

Steve: Yeah, I know. Those pesky Anonymous guys. Then they go into how they're going to pull this off, which is what I wanted to talk about.

Leo: Well, it's interesting because in fact we have identified this vulnerability before as a significant vulnerability on the Internet.

Steve: Yes. Now, the beauty, as we all know, of the Internet is that it isn't located in one place. It is inherently individual servers linked to the users through this completely heterogeneous network of interconnected links and routers, where the routers know how to send the traffic in both directions through a series of hops between any two points. So literally it's just a huge grid of interconnectivity, no single central location. The one part of the 'Net which is arguably centralized, in a sense, is the DNS root servers.

DNS, and we've talked about this, too, is inherently a hierarchy. There are the root servers which are the one place that other DNS servers can turn, or users can turn, if they want to sort of start looking for a website. There needs to be an anchor. And so these 13 root servers, which are named a.root-servers.net through m.root-servers.net. So in short, they're known as the A through M root servers. There's 13 letters, A, B, C, D, E and so forth, to M, that's 13. 13 was chosen just due to some technical record size limitations. So there's no reason that, like, 13, and it's not 12 or 14 or something, or more, it's just that that's how many conveniently fit technically in terms of the size of their name in the record.

So the argument is that, or the reason the DNS name servers come to people's attention, the reason, if this is a legitimate posting, and I should mention that this Operation Blackout, if you search Twitter for # - I had it written down here. I'm not seeing it in front of me. Oh, wait, it's in the Pastebin posting, so I could find it easily there. Oh, yeah, it's #opGlobalBlackout, not surprisingly. You can find a lot of dialogue which has occurred recently about this.

And as you and I were saying at the top of the show, Leo, the problem with Anonymous being anonymous is that they're anonymous. And so someone can post something saying that they're from Anonymous, and then Anonymous can say, no, that's not us, that's somebody else. So there's some argument about whether this is bogus or not. And I'm taking the position, well, whether it is or not, March 31st will be interesting to see if anything happens. And it may well be that nothing will happen.

Why can 13 servers withstand a big attack? The No. 1 reason that 13 DNS servers aren't going to be affected is that there aren't actually 13 root DNS servers. There are 13 IP addresses. And that's very different than 13 servers because, for example, just one of the servers, I happen to know that the "I" server, i.root-servers.net, exists in 25 different countries. And this is achieved through something known as "anycast," which is not a technology we've talked about yet.

We've talked about "multicast," which is a way of having an IP address sent simultaneously to many different recipients, or another way of thinking about it is it's a way of many recipients all asking for the same content, and instead of the server having to individually send it to individual IP recipients, they can send it to a multicast IP address, and it's automatically routed to many different locations.

Anycast is different than that. The way anycast works is that individual routers, and we've talked about, we already know routers are the way the Internet routes traffic, spread all over the globe. Individual routers have an ability to send traffic based on an IP address to the closest matching server. Now, we're used to thinking of IP addresses as being unique. So, Leo, you've got an IP address for TWiT.tv, and that's a server located in a location.

Leo: It's like a phone number. Everybody has to have a unique phone number, or you'd have collisions. Every server has to have a unique IP address.

Steve: Right. However, the technology, for example, that content delivery networks are using, is anycast, the idea being that, if you're a content delivery network, you'd like to be able to have servers stationed on both coasts, East Coast and West Coast, maybe in the middle; on other continents on the globe; and you'd want the same URL being used by different people anywhere in the world to somehow find the content closest to you.

Leo: And we use that, as well, for content delivery. So when you watch our stream or you listen, Cachefly or Ustream or Justin.tv, all the various providers we use, almost all use CENs.

Steve: Exactly. And so do the root servers. There are not 13 root servers, there are hundreds of actual physical root servers scattered all over the world. So the DNS root is actually much stronger than people tend to believe. Just, as I said, just the "I" IP address, what's funny is that the root servers are the one thing, if you think about it, you cannot access by name. That is, yes, they have names, i.root-servers.net. But you can't use that to access them because they're the root of DNS. So the one thing that you ultimately need is their IP addresses, if you have to go all the way back up the hierarchy to them in order to start looking up, for example, the address of a COM server, then the address of a, like, GRC.com, and then to actually get our IP address.

So the "I" root server is 192.36.148.17. That's set in stone. Those 13 IPs are never going to change because they are hardwired into the Internet all over the place. But the actual servers behind those IPs are free to come and go as they please. Right now there's hundreds. It would be easy over time - there's probably even more, I mean, like many, many hundreds. They're easy to set up. You create one. You put it in place. You essentially broadcast your IP to a router. And then, if you're closer to that router than another server at the same IP, the anycast technology - actually it's BGP, the Border Gateway Protocol, that routers use for communicating their routing tables - the router will go, oh, I've got somebody closer. And so anybody, any traffic coming through that router to that common IP will end up going to the shortest server.

So what this means is that an attack against the DNS root, even if the attack were - first of all, it would have to be attacking all 13 IPs. It would have to also be attacking all of these hundreds of actual physical servers hiding behind those IPs. And the only way to do that is to be attacking from all possible locations in the world because the only way to get to those physical servers is to be physically close to those physical servers.

Now, I have to say we don't know, no one knows truly whether it's possible to hold them all offline. You have to attack them. You have to flood them so much that they're not able to deliver results. Arguably, you have to flood the links coming into them, or maybe

the routers on those links, because it might be that the routers are less capable of handling a flood of small queries than the servers themselves are. We really don't know. But remember, caching is another aspect of DNS that DNS absolutely relies on. These root servers are actually not very loaded most of the time because all they're being asked for is the IPs of the second-level domains, the so-called GTLDs, the .com, .net and so forth servers. So those records in the root servers have multiday expirations, maybe even longer than that, maybe weeks.

So the point is that the only time somebody refers to them is when the .com or .net or .org, that second-level domain record that they have, expires, causing them to need to update that. The reason those don't last forever is that allows the .com and .net and .org and so forth servers to move around if they need to, and for that cached IP address for them to ultimately expire. But it doesn't expire immediately.

So that's a huge, that's an important distinction. When you flood Visa or MasterCard, you're flooding one server, and you're holding that site offline immediately for the duration of the flood. If you were able to flood these many hundreds of root DNS servers, nothing at all would happen for a while. It would, and this is what we've talked about before, as you mentioned, Leo, it's necessary to, in order for individual and users to feel the effect, which apparently is what this - if this is Anonymous and not a bogus posting, and they're trying obviously to get end users to believe that the Internet is down, that they've taken the Internet down, the only way to have that happen is to keep all of these hundreds of actual physical DNS servers offline until the DNS cache drains.

Leo: And that's at least a day or two.

Steve: Oh, many days, probably, yes. For example, when I have needed my own IP address to be more agile, I've decreased the cache time so that I could change my IP address, and users and the Internet would have their records updated relatively quickly. But normally you run with multiday, sometimes seven days. Seven days is quite common. You run with a seven-day expiration because there's just no reason for it to be any shorter. You'd rather people got to your website quicker because, remember, if you have to do a DNS lookup, that delays access to your website until your browser is able to perform a multilevel DNS query, get the new IP, if it is new, and then access you. But if it's not going to be new, you'd rather let your DNS record last a long time because that's going to help people get to you more quickly. So there really is, there's an incentive for the cache being as long as practical. And it's typically a long time.

So it is, in the first place, we really don't know today what the effect of a high-bandwidth, high-transaction rate, globally dispersed, distributed denial of service attack against those 13 IPs would be. We know that in '07 there was a denial of service attack against the DNS root, and a few of the weaker IPs were hurt. Several of them crashed. Several of them were offline. But that was, like, four out of 13. The balance of that, the other nine, sailed through it without a single glitch because they were on strong connections with strong routers in front of them, and they themselves were strong servers. And so we have some calibration. And you could argue, because this is understood to be the one Achilles heel of the entire Internet because of the nature of this one focus point - that was in '07, so five years ago. Since then, I'm sure these things are even stronger than they were. And there's just been more growth of the Internet. There's more root servers, more widely distributed. I really do think we're probably okay.

But there is one very cool site that I will leave you with. I didn't make shortcut for it. I will create one for next week. There's a DNS monitoring page. If you're looking at the

show notes, Leo, you can see it there at the end, "On March 31st." It's cymru.com/monitoring/dnssumm. That's the page. It's a very nice real-time display that shows the health of the DNS root servers as viewed from many different locations around the globe, showing the root response time and a bunch of other information, also. So I don't think...

Leo: Pretty much green. But this would be fun to go look at on March 31st, anyway, just to see.

Steve: Yup. I don't think anything is probably going to happen. And also, why March 31st? Well, it occurs to me that's the day before...

Leo: April Fools'.

Steve: ...April Fools'. So if they were going to do it on April Fools', I mean, if they said they were going to do it on April Fools', then everyone would think - wouldn't take this very seriously. So they did it one day before. Maybe that's so that the 'Net will be down on the 1st? I don't know. But I don't even know if this is a legitimate posting. But...

Leo: Well, in fact, one of the Twitter accounts - the problem with being Anonymous is no one knows who's in charge. But there is a Twitter post from yesterday saying it's a fake operation. But who knows? You just don't know. That could be disinformation.

Steve: Yeah, it could be somebody else claiming that it's a fake.

Leo: Yeah. And by the way, this is nothing new. We're not telling them anything everybody doesn't know. There's a whole Wikipedia page devoted to how to do this, that talks a lot about it and even mentions Operation Global Blackout 2012.

Steve: It's already been updated.

Leo: Yeah.

Steve: Yes, in order to cover this.

Leo: And Boing Boing posted an article - all of this comes from the chatroom, thank you chatroom - that points out, probably shouldn't make too big a deal of this, but the people who would be using this Low Orbit Ion Cannon software to do this DDoS are in fact blasting their IP address out to the public, as well. So good luck with that.

Steve: Yeah, they talk in their posting about a so-called "reflection attack," where they would send queries to DNS servers with a spoofed IP, a spoofed source IP, so that those

DNS servers would then bounce that request to the root server. So they would be spoofing the root server IP so that the secondary server would think that it was a root server making a query, which is bizarre because root servers would never do that. And in fact...

Leo: That would be an easy thing to thwart.

Steve: Yes. All you would have to do is block root server queries from those 13 IPs, block those incoming, and then your server would never bounce the traffic back to the root. And the other thing is, it's not clear what they would be asking because, as DNS drains out, then those servers would have to query the roots in order to get the addresses to query. I mean, anyway, it's sort of convoluted, and it's chasing its tail. It is probably unlikely anything's going to happen. It's also really not clear how this makes sense. Like Anonymous is pro Internet but anti some factions of the Internet, like anti-SOPA and RIAA and so forth. So why does taking the entire 'Net down even serve their ends? It's not clear. Probably really don't...

Leo: I think it's bogus, but who knows. Certainly...

Steve: Yeah, but interesting topic for the show. Interesting to consider where we are.

Leo: Steve Gibson, I know where he is, he's at GRC.com, kids. And that's where you'll find SpinRite, the world's best hard drive maintenance and recovery utility, GRC.com. It's also where you could post questions, if you've got them, and next episode will be a Q&A episode, so this would be a good time to go to GRC.com/feedback. He also has lots of free stuff there, including 16Kb of the show and transcriptions, if you'd like to read along, and his show notes, as well. You can also get audio and video of the show from our site, TWiT.tv, and wherever greater podcasts are offered for free. Do get the subscription. That way you don't miss an episode. You can have a collection of Security Now! episodes on your system.

Let's see, what else? Oh, we do this show every Wednesday, if we aren't talking about coffee, at 11:00 a.m. Pacific, 2:00 p.m. Eastern, at TWiT.tv. Do tune in live. And if you miss it, don't worry, you can always download the show, except for the coffee stuff, at the site.

Steve: And we might be changing our schedule for March 7th because the rumors are that that will be the iPad 3 announcement.

Leo: You have been paying attention. That's right. We talked about that yesterday on MacBreak. And so we're just waiting for the invites to go out. And if they do, we'll just flop, flip-flop. The Tuesday MacBreak Weekly will be on Wednesday, and you'll be on Tuesday. But that's not next week, that's the week after, so...

Steve: Exactly.

Leo: We'll know, I think we'll know by next week. Usually Apple sends out invites about a week ahead of time. Thank you. Actually, you know what, next Wednesday is the day that Microsoft ships the consumer beta of Windows 8.

Steve: Ah.

Leo: So that might be kind of interesting.

Steve: Yeah.

Leo: We'll talk about that. Thank you, Steve Gibson. Thanks for joining us, everybody. And we'll see you next week on Security Now!.

Steve: Thanks, Leo.

Leo: Bye bye.