Listener Feedback Q&A #12

Description: Steve and Leo discuss questions asked by listeners of their 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 they have previously discussed.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-064.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-064-lq.mp3

INTRO: Netcasts you love, from people you trust. This is TWiT.

Leo Laporte: Bandwidth for Security Now! is provided by AOL Radio at AOL.com/podcasting.

This is Security Now! with Steve Gibson, Episode 64 for November 2, 2006: Your questions, Steve's answers.

Security Now! is brought to you by Astaro, makers of the Astaro Security Gateway, on the web at www.astaro.com. And by Dell. For this week's specials, visit TWiT.tv/dell.

It's time for Security Now! Episode #64. Let me think. There's something about that number. Oh, I know. That's how many squares that are on a chess board. Eight by eight. So that would be four by four by four.

Steve Gibson: By four.

Leo: By four. By four, by four. Four by four by four, which would mean, ladies and gentlemen, it's a mod 4 episode. And time...

Steve: Indeed it is.

Leo: ...for your questions and answers. Hello, Steve.

Steve: Hey, Leo.

Leo: So we had some fun with MojoPac last week, and I hope people enjoyed that one. I'm sure we have questions about that and a lot more. Should we just get right to them, or do you have anything you want to catch up with?

Steve: Nope, I think we're ready to go with – we got a dozen listener Q&A. And I have to say, Leo, as I'm reading through these postings, I did want to remind people. People seem to be finding it anyway, but at the bottom of the Security Now! page at GRC, which is just GRC.com/securitynow, the bottom of the page is a web form that anyone can use to send me questions, which I receive and, you know, do my best to read through. And...

Leo: We should emphasize, you do not answer individual questions.

Steve: I get so many of them that I can't. But every time I'm preparing one of these mod 4 shows I find myself thinking, wow, you know, I really want the ability to answer more of these. So one of these days we may actually, you know, as I have threatened to before, switch maybe to mod 2 in order to double the frequency of doing Q&A. I really think people like them because, you know, one episode then covers many different topics. It allows us to talk about, you know, small things rather than, you know, just single big things. And they're just, you know, we've got such an active, really interested audience that it's really cool.

Leo: Well, let's kick things right off with question number one from Michael Klein, Fairfield, California, who says: I have a question for you. It was on your topic a couple of episodes ago, Proxies. We only talked about one kind of proxies. But: The corporate network I work on requires using a proxy to access the Internet. Obviously this allows them to track everything employees are looking at, even if the data might be sensitive personal information. Now, I'm using SSL connections. You indicate that using an SSL connection prevents the use of transparent proxy by ISPs. I'm wondering if an SSL connection prevents snooping by a corporate proxy, or can a corporate proxy intercept the SSL connection. For instance, could a proxy set up two SSL connections – oh, he's smart, this guy's smart...

Steve: Yeah.

Leo: ...a client proxy connection and proxy web server connection, decrypting the information, scanning the content, reencrypting the information and sending it along – kind of a man-in-the-middle attack. I think they do that, don't they?

Steve: Uh-huh. And he says: Can a proxy fake a certificate to make you think you have a completely secure connection between a client and web server when you really don't? Now, we've talked about this a couple times. And in fact, this is a great opportunity for me to clarify something that I wasn't clear about relative to certificate verification.

So first of all, Michael is exactly right that there are corporate proxy servers which do intercept; and, exactly as he said, the web server sets up an SSL connection to the proxy. The proxy sets up an SSL connection to the target web server. But in between there it decrypts, can scan, and typically does scan for content and then reencrypts on the way out. However – and again, he's on the ball because what this requires is some sort of a certificate game being played. So in fact proxies do perform – corporate proxies can be set up to perform exactly this kind of interception. However, in every case, the browser must be complicit in this. That is to say, what happens is the proxy will have an SSL certificate, and it is necessary that the matching

certificate be preinstalled in every corporate browser. Now, again, that does happen and is happening to people who may or may not be aware of it. But what then happens is that when the user attempts to use an SSL connection, for example maybe to, like, go to Gmail and use Gmail's web browsing mail interface in order to have a truly secure and private interaction that is going to truly bypass the corporate proxy, the proxy will itself return its own certificate.

Basically what it does is it creates a certificate for that remote website on the fly, which it signs. Remember that in the past we've talked about certificate and certificate signing, and listeners can go back and look down at our crypto series to get the details of that. But basically it creates on the fly a security certificate from that website. Now, normally, for example, Google would have a certificate that was signed by Equifax or by VeriSign or one of the major certificate signers. And everyone's browser has those large companies – Equifax and VeriSign and Microsoft and so forth – they have those certificates already installed in their browser, so that they're able to validate the signature of the web server's signed certificate. So what happens in the corporate scenario is that the corporate proxy creates basically a fake certificate on the fly, which it signs. It itself, that corporate proxy signs it. That's what it returns to the browser. The browser then looks in its cache of certificates, and it will find Equifax, VeriSign, Microsoft, and Ajax Company, or whatever the company's name is, so and it'll check to see that the signature of the remote server certificate is properly signed, which it will be. So no alert will be presented to the user. The user won't know that anything untoward has happened. And this ends up being a seamless connection.

Leo: Can't you, though, look at the certificate and see all of that going on?

Steve: Ah. Now, that is exactly the point I wanted to make, and what I wanted to make very clear about how you can bust this happening. And that is, when you bring up a secure page, right-click on the page and look at properties of the page. There's a View Certificate option. What you want to look at is what's called basically the Signing Chain. And if you look at that, it'll normally say, like, Google.com, and then there may be an intermediate authority, but then there's a root authority which will be VeriSign or Equifax or Microsoft or whomever signed it. Well, in this case, that's the giveaway. It'll be Ajax, Inc., or whatever the name was given to your certificate, probably the company's name. So you'll know immediately that the certificate you have received has actually come, not from the actual server signed by a true universal authority, but it was signed by your own local proxy server, meaning that the connection terminates at the proxy server where it's being decrypted for whatever purposes.

Leo: Although I have to say, somebody at work sees, you know, the company name, it's probably not – it's not a red flag. They're going to go, oh, yeah, oh, no, I'm going through the company. But now you know that's a giveaway the company may be intercepting the traffic.

Steve: Well, actually you know that they are definitely intercepting the traffic. You don't know what they're doing with that decrypted phase because, if the certificate your browser has received from what it thinks is the website is, in fact, signed by the company, then the data has to be decrypted as a result of it traveling through the SSL connection and coming out at the proxy server, where it's reencrypted. Now, it may be that nothing is happening except the company is just adamant about proxying everything, and so they've taken these measures. But I do know that this is being done by some proxy servers and companies that want this level of control. So it's always possible to detect it in that fashion.

Leo: There's no way they could do it without having a certificate that said something other than the site you're going to.

Steve: No. What you would – if they tried to do that, you would get a popup dialogue warning you...

Leo: Saying it doesn't match, ah. I've seen that error. Okay.

Steve: Yes, yes.

Leo: Okay. It says the certificate doesn't match, what do you want to do?

Steve: Yup.

Leo: And you can accept it temporarily or forever or whatever. Jon, lurking somewhere in New York, writes: Could you please talk about SSL? Michael was talking about SSL. I really don't understand what it is, what it means, what it does. I mean, I sort of know what SSL connections do, but I don't know how they actually work. That's a great question.

Steve: Well, it's so great that I'm going to give that an entire episode. I thought I'd answer it really quickly, sort of in an overview mode. But and I thought we had really covered it. So I looked back at our work after our crypto series, and I never really addressed the mechanism of SSL. Its origins were with Netscape. Those were the guys that came up with SSL 1. And it turns out there were a couple sort of boundary condition problems with it that were fixed. But basically the guys at Netscape did it right. And now SSL, also known as TLS, which is Transport Layer Security, I mean, it's such a great solution that it's worth really just stopping and explaining in detail how that works with certificates and crypto and everything. And we do have all the background, having really covered cryptography very well. So I want to do that.

But, you know, the short version of Jon's question, or the short answer to Jon's question, is that it's a means for allowing true authentication, if you choose true authentication, and very strong, industrial-strength encryption of data through a channel, such that a man-in-the-middle attack cannot be performed if you are securely authenticating the endpoints, which SSL does provide, thanks to certificates. And somebody passively sniffing the traffic has no ability in, you know, the lifetime of the universe sort of timeframe to decrypt that traffic. So it is a, you know, it is proof against anyone intercepting the traffic. It allows you to authenticate, that is, to prove the identity of the endpoint, given that, for example, as we were just talking about with Michael's question, that you check the certificates. That's certainly important. But it provides these mechanisms. And it's just a – it's a terrific means for moving data across the Internet, solving problems of security and privacy. And we're going to give it a whole answer in its own episode.

Leo: Good. I mean, of course people have seen it when they go to their bank; you said Gmail. My email provider uses it. It's basically relatively free and a good thing to use. I mean, you just have to get a certificate.

Steve: Well, and anytime you see https...

Leo: That's secure.

Steve: The "s" stands for "secure," which is SSL.

Leo: Great, yeah. Tom Krauska of St. Louis, celebrating his World Series, has been having trouble – actually, as we record this, we don't know how St. Louis has done in the World Series, but I'm sure it's been wonderful – has been having trouble sending email to a few specific people. He asks, could you guys talk about how to get around Port 25 blocking? I have email I send to my sister and a friend, but it returns to me with this line: Connected to (some IP address), but sender was rejected. Remote host said: 550 (another IP address) blocked by ldap:ou=fblmx, dc=Comcast, dc=net. What?

Steve: Yup. We've never talked about what are called "real-time blacklists." And so Tom's issue is not that he's got a problem with Port 25.

Leo: Oh, there's more. I didn't see, there's another page. BL004 Blocked for spam. I left that out. He said, I've been told Comcast is blocking Port 25. My ISP says only support email on Port 25. This sounds like a gunfight at the OK Corral. Well, so he's not a Comcast customer. He's trying to send to Comcast. Is that what's going on?

Steve: Well, whichever way, because it isn't clear because we don't have his IPs.

Leo: Right.

Steve: But essentially he doesn't have a Port 25 problem.

Leo: Right.

Steve: He's got a false spam triggering problem.

Leo: Right.

Steve: Real-time blacklists are essentially lists of known or believed spammers. So what's happening is, he's able to send email apparently just in general to lots of people. But his sister and this friend are unable to receive his mail. So what's happening is that the SMTP server at their end, that is, their ISP's SMTP server is doing everything it can to block spam. So it's accessing these real-time blacklists and comparing the IP of his SMTP server against the list and finding that that IP or that range of IPs is believed to be a sender of spam.

Leo: He's been blacklisted.

Steve: He's, well, he's been blacklisted. But not his IP, actually his ISP. Because...

Leo: Most likely his ISP, right, yeah.

Steve: Exactly. So...

Leo: Oh, it could be him, it could be him, but it's most likely his ISP.

Steve: Well, it can't be him because the connection is coming from his SMTP server.

Leo: Oh, that's right, that's right.

Steve: Not actually from him directly. So...

Leo: Unless he's running his own mail server, but that seems unlikely.

Steve: Exactly. So the way – so, you know, even for a guy who's innocent, I mean, you know, Tom could be completely innocent. But what's happened is other customers of his ISP are infected with spam-sending bots and zombies. These things are – so other people's machines are infected. They're depositing spam on the ISP server, which is the ISP's SMTP server, which is in turn relaying it, as spam servers do, all over the 'Net. People are receiving this spam, backtracking it to this ISP's SMTP server, and blacklisting it.

Leo: You're being kind. You're being kind. Because unfortunately the problem with real-time black holes is you can get black-holed for – blacklisted for all sorts of things, maybe not even sending spam.

Steve: Well, and that's the problem, yes, is there's a high level of false positives in all of this. I mean, again, it's...

Leo: Most of these – as we know, most of these zombies don't go through the ISP's mail server, they do their own mail server. That's why these ISPs block Port 25.

Steve: Right.

Leo: You would get caught very quickly if you tried to forward mail through your ISP's mail server. So it may be that the ISP had, you know, did some web hosting for somebody who was a spammer. There's lots of ways. I've had addresses get on these black hole, blacklists, and it's very hard to get them off.

Steve: Yup. You just have to wait.

Leo: Yup. MAPS/ORBS, he had an rbl, what was it, rblmx something? So what is that ldap? That's just a lookup of that?

Steve: Yes.

Leo: Okay.

Steve: Yeah.

Leo: Rblmx, blocked by ldap:ou=rblmx. So it could be his ISP was at fault, or somebody on the ISP is at fault, or it could be the black holes, blacklists at fault. But either way he's not going to get email through. Is there any recourse, anything he can do?

Steve: Well, okay, here's the problem. As you said, ISPs are blocking Port 25. So what that means is that the only thing he could do would be to find some other email server...

Leo: So use Gmail, use Yahoo!, use Hotmail.

Steve: Exactly. If he sent, you know, basically he's not able to use his ISP's email. His ISP is blocking his direct connection to other remote servers, so he can't do that in order to avoid his blacklisted SMTP server. So the only thing he can do is use a complete third-party email solution.

Leo: I think if you sign up for a Gmail account you can use their SMTP server for outgoing mail. And that would be a free way to do it. I believe you can.

Steve: I know that you can receive using POP. You're able to pull email...

Leo: Yeah, I think you could send out, yeah.

Steve: I don't think you can.

Leo: You can't? Oh, okay.

Steve: I don't, I mean, I have looked for doing that, and I have not seen a way to do it.

Leo: FastMail will allow you to do that. That's a paid, of course, service. And I bet you if you paid for Yahoo! Mail, you'd be able to use an SMTP server from them.

Steve: Well, and again, that would open them up to being a spam relay. So...

Leo: Well, but if they start seeing a lot of traffic coming from one person, then they just cut them off; right?

Steve: Can you pay for – oh, no. You can't pay for Gmail. You were talking about Yahoo!.

Leo: No, the Yahoo! Mail, yeah.

Steve: Okay, yeah.

Leo: Jason Bennett in Chesterland, Ohio is a bit worried. He asks: With all of the IE issues that always seem to come up, I thought of a question. Is there the same vulnerability issue if you use the IETab plug-in for Firefox? That's a neat plug-in that lets you continue to work in Firefox, but actually open a site in IE within Firefox.

Steve: Yup. And of course that's a problem. When you open the tab, you are basically running the Windows IE control. This is the same problem, for example, that – and it's the reason that Outlook Express or Microsoft Outlook has email problems, because that preview and the actual file viewer is the same viewer, it's the Internet Explorer viewer that allows the display of web pages in mail. Similarly, when you use the IE tab, they haven't obviously brought along an entire Internet Explorer version of their own. They're actually using the component ability of Windows to just take the IE Window and move it into the Firefox tab. So any security vulnerabilities that exist in IE will bite you if you use the IETab under Firefox.

Leo: Same is true anytime IE is rendering a page – in Outlook, in an application that uses IE to render HTML.

Steve: Yup.

Leo: You're always at risk. Hey, I just looked it up. I thought I remembered an article on Lifehacker. You can use Gmail as an SMTP server. SMTP.google.com. You'll have to use...

Steve: No kidding.

Leo: Yeah. You'll have to use your Google account. You have to have a Gmail account, obviously. And you'll use TLS under the secure connection. And you can use their SMTP server.

Steve: Very, very cool. I'm going to do that, Leo. I actually have an application for sending stuff through Gmail. I didn't know that would work.

Leo: Right. Well, and I remember a number of people I've talked to have hard times sending email out in various situations, like when they're on the road and so forth. So they just use Gmail. Very neat. Thank you, Google. We'll see if they keep that up now that we've mentioned it. That may be the end of that.

Phil Hendy from Swindon, that's in the U.K., says: After using ShieldsUP, my PC, behind a router, passes all tests. It's stealthy. Except for the ping. ICMP echo request. I can't find a way to disable this. And if I did manage to do this, what limitations, if any, would I face?

Steve: Yup. We get this question a lot, so I wanted to pause. And, you know, Leo, you doing that British voice reminds me, I would love our listeners to hear you switch to DJ mode.

Leo: Why would you like that, there, Steve? It's a beautiful Wednesday afternoon, a

Thursday afternoon. We're doing a Security Now! right now, 12:40 in the city. 58 degrees. So why would you want that, Steve?

Steve: I just think – that cracked me up. You...

Leo: Makes you happy.

Steve: You've done it a few times in the Call For Help studio. And I've – the first time I heard that, I said, oh, cool, you know, Leo has that whole, you know, DJ voice down.

Leo: Hey, when I worked in radio I had to talk like that. Wait a minute. I still work in radio. We – people talk like that. I don't know why.

Steve: Wow.

Leo: Anyway...

Steve: Okay.

Leo: Let's get to Phil's question...

[Talking simultaneously]

Steve: Yeah. A lot of people ask the question, so I wanted to address it. He mentions he's behind a router. So what's happening is, he's of course got a NAT router, which is inherently going to give him tons of incoming protection.

Leo: Yes.

Steve: So naturally he's going to be fully stealthed. Some routers, however, will still echo a ping. So he's going to fail ShieldsUP full stealth, or true stealth test, as I call it, because I do send a packet just to see if there's a response. What he can do is, first of all, he might look around the user interface. Some user interfaces have a very clear ability to specifically turn off echoing ICMP pings. So if you – so I would first say, look through the router's security configuration or WAN security or WAN tab to see if it already offers you the ability to turn off ICMPs.

And if it doesn't, what you can do is use the DMZ feature. We've talked about the DMZ, the so-called "demilitarized zone," which is what DMZ stands for, where unsolicited traffic is routed to a specific machine on your network. This is what people have to do, for example, if they want to run a web server or FTP or, you know, if they explicitly want to send unsolicited traffic somewhere. That's where they use the DMZ to specify which machine gets it. This is not the normal configuration because, as we know, what makes a NAT router such a good security appliance is that unsolicited traffic is simply dropped, and only traffic coming back from connections that were first established from the inside out are allowed back from the outside in.

But all you have to do is configure your router's DMZ to a nonexistent IP. For example, if in your private network behind the router you were 192.168.0.1 through 50, for example, you could just aim it at 192.168.0.222. Something, you know, high up out of the way that's not ever going to be an IP of a machine. So what happens is, that ICMP packet and any other traffic, it would, being unsolicited, come through the router; but there's nowhere for it to go inside your network, so it simply gets dropped, in the same way that the router would normally be stealthing.

Leo: That's very clever. I never thought of doing that. Now, most I've seen, almost all routers have somewhere in there a way to turn off ICMP.

Steve: Yes.

Leo: But I have to say, I bet you you could make a lot of money by just going around to people's houses, Steve, and looking over their shoulder while they navigated through. I'm just looking through my D-Link menu, and there are so many places it could be hiding. And I'm just...

Steve: Yes. Routers have gotten very sophisticated.

Leo: And there's no standard. There's no way of...

Steve: One of – right. I'm sorry to interrupt. One other thing that is worth mentioning is you might also upgrade your firmware. People don't generally upgrade their router firmware, you know, unless they have some specific need to. But router manufacturers are more or less continually fixing problems and adding features. So it might well be that when you bought your router it did not offer the ability to bock ICMP echo requests from going back. But if you just brought it to a current firmware build, you might find that that feature had been added.

Leo: All right. Moving along to another question here. From U.K. to Minnesota we go. A listener named Brian has been pondering virtual machine technology. We did a couple of episodes on that. He writes – more than two.

Steve: Yeah, we covered it pretty – we beat it to death.

Leo: Yeah. He says: I've been hearing that only Windows Vista Ultimate Edition – oh, this is the new license agreement thing – will be able to run in a virtual machine. I guess I understood that there is no way the OS would know it was running in a virtual machine. Can Microsoft enforce this; and, if so, how? This is actually a conversation Paul and I had on Windows Weekly when we were talking about these new license agreements. And he was of the opinion it's unenforceable.

Steve: Independent of the legal enforceability, it is certainly possible for Microsoft to detect this. We talked about this during the episode on Blue Pill. Because remember that the idea was that you could run some software that would switch your system into virtual machine mode on the fly, literally without a reboot. Suddenly, a hypervisor, as it's called, would be loaded and run. And, I mean, literally your screen doesn't even blink. Nothing happens. It's just suddenly

been subverted.

Now, the question was that, you know, whether this is a detectable thing because with modern virtualization technology everything can be captured. Because the chip itself, the processor itself, so ports virtualization at the hardware level, there isn't a way sort of broad brush-wise to detect it. Except it's so difficult to do a perfect job of virtualization. So, for example, you might be able to virtualize a machine more easily that was not on the 'Net. But, for example, time is very difficult to virtualize, that is the passage time.

I mean, you could argue that, okay, well, if we knew that certain instructions that were going to take longer when the machine was running in a VM mode, their timing would be different. And it is the case that instructions are – some instructions that have to go through this extra hardware virtualization take way longer than when they're not. So you would imagine that software could simply time the execution of the instruction. But the timer itself can be manipulated by the virtualization layer so that the time appears to be the same. The problem is, if you have access to the 'Net, you are able to access time servers. Now the job would require intercepting them, too, and faking them.

So I guess my point is, if Microsoft really wants to catch people running Vista in VMs, and you have to imagine Microsoft does really want to catch them, given this growing propensity that we're seeing on Microsoft's part of, you know, bolting down the licensing and enforcing it with technology, I'd have to imagine that Microsoft would be able to penetrate the virtualization layer and detect that something is going on that they don't want to allow.

Leo: Bottom line is it is detectable, it's just hard to do. You'd have to be fairly motivated to do it, to make the effort to do it.

Steve: Do we know what Ultimate Edition is going to cost? Is there any pricing structure?

Leo: Yeah, I think it's 300 bucks. It's very expensive. That's if you buy it, you know, straight out of the box. There are upgrades and so forth.

Steve: Right.

Leo: So, yeah, I think that – I don't know. I mean, they know, for instance, if you move a virtual machine to another machine, they ask for activation again now. So they're getting fairly smart about this. I would gather – I would guess you're probably right, they are motivated. They've certainly written into the – well, I guess their premise is, if you're a home user, you're buying Windows Vista Basic, what are you doing a virtual machine thing for?

Steve: Exactly.

Leo: But as I pointed out to Paul at the time, well, I think more and more users are going to start using virtual machines, maybe even home users who want security. So, let's see, here. Daryl from Kansas says – oh, it's a MojoPac question, our last episode – I tried MojoPac and moka5 USB solutions, both of which you mentioned on Security Now!. Then I thought, hey, how do I know I'm not installing a rootkit, spyware or whatever by installing these? Question: Do you check out that stuff before mentioning these products on your show? I went to McAfee SiteAdvisor, keyed in moka5.com, it came back with a green

check, meaning we tested this site and didn't find any significant problems. I did the same for MojoPac and received the message "No results found for mojopac.com." So I submitted it for review. Do you do that?

Steve: Yes.

Leo: Of course you do.

Steve: I will never casually mention a product, you know, from some questionable site. And in fact, people who have listened to all of our podcasts will hear me, you know, very carefully say, when I do mention, you know, an underground site or a site which is known to be a hacker or cracker site. I mean, I go to great pains to make sure people understand. There's just no way that, you know, an upfront company like MojoPac or moka5, where, you know, I go over their site, I look at their background, I figure out what's going on. I mean, I do – I'm not saying that I'm, you know, checking in detail everything that these companies are doing. But I'm certainly looking to verify that these are, you know, well-known, reputable companies, you know, financed with venture capital, or founded by ex-Stanford professors in the case of moka5. And, I mean, so I'll be very careful when I talk about things that could have a negative consequence, and caution people that, you know, that's the kind of URL we're talking about.

Leo: This is the guy who discovered spyware, Daryl. C'mon, he's not going to put spyware on your system. But we should point out that, you know, there are lots of nefarious, sneaky people who are trying all this time to do this. So mistakes could happen. But Steve, I know, is assiduous in protecting you and would never – checks everything.

Steve: Yup.

Leo: Never let anything go by. Agent Smith – obviously not his real name – hiding somewhere in deepest Australia writes: Well, we experienced our first really serious DDoS attack a few days ago, and it was an exciting experience. Exciting in the way I imagine being shot is, shot at is exciting. There were network engineers doing what they could while senior management's attitude was, if not explicitly stated, "Get us back on the air now." Anyway, if memory serves, you and Leo covered DDoS in your "Internet Weaponry" episode a few months back. Steve knows more about DDoS than he even lets on.

Steve: Yeah.

Leo: Even then, have there been any new countermeasures developed to help deal with – DDoS, they're Distributed Denial of Service attacks where a bad guy gets a bunch of machines, thousands, maybe more, attacking you all at the same time. I'm guessing that the only way to reduce the impact is to install big pipes, redundant network hardware, and not to piss off the various attackers. But I thought I'd ask. Apologies for not giving my real name, mate. Management's a little spooked by the experience and want to keep it as quiet as possible. Rightly so.

Steve: There are – this is a great question. And of course I sort of keep my eye on what's going on with Denial of Service.

Leo: Because you've been DDoS'd.

Steve: Have been in the past and, you know, have had lots of interesting experiences, you know, talking to ISPs and to other users. The bottom line is that bot networks now, these remote-controlled trojan, zombie, whatever you want to call them, bot networks are so large that dealing with a denial of service attack that is sufficiently big is really a problem. It's important to note, for example, though, that the size of the attack is very important. In the old days, when I was running the GRC servers here with me, only at the other end of two T1s, I had 1.54 megabits twice, so a little over 3 megabits of total bandwidth. So, you know, three computers, or maybe ten, but certainly not hundreds or thousands, you know, a small handful of machines pounding on me could easily saturate my pipes at that point and, you know, take GRC off the 'Net because the servers at the far end of the T1s would be filled – the connection would only be delivering a high percentage of garbage and not any of the good content. Now, there are companies that are in the business of protecting people from big denial of service attacks. And again, it's all a matter of scale.

Leo: What do they do? Just they throw pipe, throw bandwidth at it?

Steve: Yes, exactly as Agent Smith here has said. They have really, really big pipes, and they set up some proxies so that the proxy responds to the TCP SYN request to establish a connection. And then if it's a spoofed SYN, the proxy will respond to it. But there isn't a machine at the spoofed IP address so that their SYN/ACK going back from that big pipe proxy just flies off to Never Never Land and is never heard from again, only real people who send a SYN because they want to establish a TCP connection to the server. They respond to the SYN/ACK; the proxy then turns around and opens a connection to the actual web server that is cowering behind this big pipe proxy server and receives the requests. So the attack flood is completely blocked that way.

But these big pipes, meanwhile, are being flooded with stuff. And if the pipes are not big enough, or even if the proxy servers that are attempting to respond to all this SYN traffic are not fast enough, you'll still end up limping along a little bit. However, this is what the gambling sites that have been under such trouble with denial of service attacks, and other sites that really have to stay on the 'Net in order for their well-being to be preserved, these are the kinds of services that they've had to move to.

And notice that these big pipes are expensive, even when they're not being used. So what's happened is it's changed the cost structure of Internet connectivity. You can get very inexpensive Internet connectivity that isn't protected like this; and if you're not a target and you don't piss off anybody, as Agent Smith said, you're probably okay. If you're a high-value target and it's really important that you stay on the 'Net, then what you end up doing is paying a lot more for your Internet connectivity because you're having to finance your shared piece of the infrastructure to deal with attacks when they happen, even if they don't.

So, I mean, it is something that has raised the cost of bandwidth for people who are high-value targets, and but there is really no other solution. There isn't anything, I mean, even theoretically possible that a small company can do to prevent themselves from just being taken off the 'Net during a significant denial of service attack. Again, they could increase their bandwidth, increase the, quote, the "size of their pipes" in order to withstand larger attacks. But attacks now are so big because there are literally tens of thousands, even hundreds of thousands of bots in these networks that they're just – they're able to produce so much bandwidth, just unplugging your equipment and taking a vacation is the recommended course of action.

Leo: Wow. So another way people do this is by changing to a different IP address. And but then, if the attack is aimed at not the IP address but the name, that's a problem.

Steve: Well, Leo, and I don't have to tell you about DNS propagation problem.

Leo: Yeah, and that's the other thing, sure. Wow.

Steve: You know, you've had those for a few weeks when you changed servers. It takes a long time for...

Leo: It does.

Steve: ...the 'Net to catch up when you change your IP.

Leo: We've talked in the past about various schemes for kind of defeating denial of service attacks. Did none of those really pan out?

Steve: Well, I mean, theoretically that much bandwidth has to go somewhere.

Leo: You just can't– if somebody's flooding your pipe, it's flooded.

Steve: Yup, it's over.

Leo: There's nothing you can do. Brian Scallan asks from London: The new Firefox 2.0 browser privacy settings allow – by the way, it just came out this week, or actually last week – allow or disallow cookies without distinguishing between first-party and third-party cookies. Mm-mm-mm, like Internet Explorer. There is an option to explicitly allow or disallow cookies from user-specified sites. But even then it doesn't distinguish between the different types. It would seem to be an all-or-nothing approach to cookie handling. Is it less safe as a result? Wow, that's bad news.

Steve: It is so bad, Leo. I'm so...

Leo: I can't believe that.

Steve: I am so disappointed. They removed – under all prior versions of Firefox up to 2.0, there was a nice setting in the UI that said "Allow cookies from the target site only," and you could turn that off, and you got exactly what I recommend and what we've talked about, and we will be in the future because, you know, I've invested heavily in third-party cookie awareness on GRC for a bunch of stuff that hasn't yet gone public. And I know what a percentage of users have third-party cookies on.

Leo: What is – would you describe what a third-party cookie is?

Steve: Well, the idea is that a cookie is a token which a site you're visiting gives your browser, which then it returns for successive requests. The idea is that – and this again is something that originated, as did SSL, originated from Netscape. When a browser makes a connection and requests a page, that's a – we'll call it a transaction, a web transaction. It asks for the page, then it asks for all the GIFs and JPGs that are there and other assets in order to populate the page. Then the user sits there and looks at it and decides what they want to do, and then may click a link elsewhere on that site or somewhere else on the Internet.

Well, if they go to a different page on that site, the browser makes a new request, asks for the page and all the other goodies that the page contains, and then the user looks at it again. The point is, there is no way to link those two pages. That is, there's no way to know that the same user who looked at the first page then looked at the second page because they are completely separate transactions. Now, if we were in a proxy-free world, the server might look at the IP and notice, oh, look, we got a request from this user's IP a minute ago, so it's the same guy. Except that we know ISPs – we just finished talking about ISP caching proxies – ISP proxies will have the IP which is coming to the server. So the server cannot use the apparent client IP in order to disambiguate all of the traffic coming to it.

So what Netscape did was they created this thing called a "cookie," which is part of the headers, it's part of the request and the reply headers at the beginning of the query or the response from the browser or from the server. The idea is that, when the client browser makes a request, it checks to see if it has a cookie, as it's called. A cookie is just – it's called an "opaque token," meaning that it's some blob of data that has no particular meaning to the browser, but it does mean something to the server which gave the browser the cookie. Normally it's, you know, some – a little blob of crypto stuff or a serial number that's been encrypted or something. The idea being that the browser, if it's got a cookie – if I'm at Google.com and the browser's about to make a request, it looks to see if it has a Google.com cookie. And if so, it adds that to the headers in the request. That allows Google to check its database and see who it is that this user is. For example, it allows you to stay logged in to Gmail over time. It allows, you know, this is the mechanism that allows eBay to know, once you log in, who you are as you move along.

There are other approaches that involve putting this sort of information up in the URL, but cookies are a very nice means for doing it. And, I mean, and they work really well. The problem is, and this is something that Netscape sort of considered but didn't get worried enough about, is that all assets that are populating the page are cookie-enabled, not just the main query for the page, but, for example, the images, as I mentioned.

Leo: Banner ads, images, all of that stuff.

Steve: Well, exactly. Banner ads are the trick. Because companies like DoubleClick, major ad services, realized that, hey, we can give users a cookie, too, when they ask for an image on our page, you know. And that's a unique token that follows the user around the 'Net. And that's the problem, is these third-party cookies, as they're called. A first-party cookie is a cookie between you and the site you're visiting. A third-party cookie is a cookie which your browser gives when it asks for an asset, not from the site you're visiting, but for example from an advertising server's site.

Leo: Now, why does it matter if DoubleClick gives me a cookie when I go to the CNET page?

Steve: Well, what matters is that there is – well, essentially, that cookie you get when you go to the CNET page is returned to DoubleClick if they're an advertiser on MSNBC.

Leo: Uh-oh.

Steve: Or an advertiser on Microsoft. Or on...

Leo: So they can follow me around. CNET can't follow me around. MSNBC can't follow me around. But the advertiser that's common to all those sites absolutely can.

Steve: Yes, it's third...

Leo: They can know where I've been.

Steve: It's third-party cookie tracking, and it is absolutely done everywhere possible, and essentially to compile a dossier, a profile of all the sites you visit. Now...

Leo: Now, you may say, well, that's harmless. But if you give one site some personal information and they can track that and combine that with the cookie, suddenly they're starting to really build quite a bit of information up about you.

Steve: Yes. In fact, there are sites that have been set up – there were sites that offered, like, free prizes if you just, you know, it was like, here, here's a chance to win a free prize, fill out this form. These sites are operated by the advertising companies. They contain ads from the advertising company. Thus they are supplying that advertiser cookie, and they provide the data you fill out – well, I mean, it is the advert- it's a front put up by the advertiser company to get your name, address, physical address, phone number, email address, and essentially to acquire non-anonymous information from users.

Leo: Oh, interesting.

Steve: Then they're able to sell this – and they then merge that with everything else they know about you. They know if you went to some search site and spent three weeks searching about HIV/AIDS. Because another problem is the so-called "referrer." And we'll talk about this because, I mean, I've got a real issue with third-party cookies. The referrer in searches contains the information that you search about. It turns out that's provided in the query to the advertiser. So advertisers not only know all the sites you go to, they know what things you search for when you use search engines.

Leo: And as we have seen in the past, you compile enough information like that, you know everything about somebody. I am – I'm just – because I couldn't believe this, just downloaded and installed Firefox 2.0, and he's absolutely right. They're – the only cookie management is to turn it off or leave it on. They've completely gutted what used to be a pretty good cookie management tool.

Steve: Yes.

Leo: Now you have one checkbox that says "Accept cookies from sites," yes or no. And then you can have exceptions. So you probably would want to say, you know, put DoubleClick.net and always, you know, if you could add that list, at least that would let you block the major guys. But it's not. It's not a good alternative.

Steve: The good news is, this is configurable in Firefox 2.0, but not through the UI. You need to use the configuration file. And it is possible to still say, I only want to accept first-party cookies; I want to deny third-party cookies.

Leo: So they just don't expose it in their regular interface.

Steve: I have no idea why they took it off the UI, but yes.

Leo: You do about – about:config will show you all of these settings.

Steve: Exactly.

Leo: And do you know what the setting is that we have to change?

Steve: Yes. It's network.cookies.something.

Leo: All right. Let me look for it.

Steve: We're going to talk about it, and my site will be talking about it. I don't have it up yet, so don't go looking for it. But they must have just run out of UI space. I mean, they must have had other stuff they thought was more important, they didn't want to – you know, just feels like why would you remove something like this? And they may have figured that people don't care or don't know. Well, it's certainly the case that people don't know. I believe, you know...

Leo: They care if they find out.

Steve: If they understand. And so I've invested a lot of time and engineering over at GRC, and it'll be another major feature of GRC. In fact, it's – I did the menuing system because cookies was something I wanted to get done, as well as the long-awaited OpenVPN documentation, both of which are waiting for me to get the menus online, which is the next thing I'm going to do when I get back to GRC. And anyway, but...

Leo: I'm looking at the – I'll tell you, there's no obvious place. I mean, I see a lot of network cookies, always accept session cookies, network cookie – there's one that's called network cookie, cookie behavior.

Steve: That's the one. You set that. I think it's normally zero...

Leo: It's zero.

Steve: ...and you set it to one. And it only allows first-party cookies.

Leo: You wouldn't even know to set that. It doesn't say anything about first-party/third-party. So, again, you could surf to about:config and change network cookie – network.cookie.cookie behavior to 1 by double-clicking it, and that's going to do it? Wow.

Steve: Yes. And if you do a Google search for Firefox 2.0 third-party cookies, you'll find a bunch of controversy about this, and dialogue. And these instructions are available there.

Leo: That's very disheartening, that they would take such a big step backwards.

Steve: Yup.

Leo: All right. I'm sorry. I'm just very upset. Gary Swann sends us a great question from England: If my ISP uses a transparent proxy – transparent proxy – how do I know if I'm getting the latest web page? Even if I reload the page, surely I get the page from the proxy cache instead of the latest page from the web server. Is that true?

Steve: He's exactly right.

Leo: Wow.

Steve: If your ISP uses a transparent proxy, there's no way you can get around it. You know, all of your web access is going through your ISP. They're filtering it. There's nothing you can do. And there's no way around it. I mean, if the cache is not set up correctly, it will not be – it will be keeping a copy and feeding that back to you. There just – there isn't a way to avoid that. Hopefully the cache is configured so that – and the web server on the far end is configured – so that if pages need to kept fresher than the web server is serving content that says "Do not cache this," or "This expires in five minutes," or whatever period of time. And so, again, it requires that the rest of the system be functioning correctly. But if you've got a cache in the way which is misbehaving, there's no way to get around it.

Leo: Wow. Wow. That's just the nature of caching.

Steve: Now, you could do something like, you know, go to a third-party service like Anonymizer or a TOR proxy that we will be talking about. So you could come up with another way to circumvent the transparent cache by not routing your web traffic through your ISP at all. And so, you know, there are, like, you know, ways like that. But it, you know, requires a lot of work.

Leo: Right. Monty Janak of Needville, Texas was wondering: If a virus or worm edits my DNS config, could I get sent to a malicious DNS server website on the Internet? This seems

like a security issue similar to editing the hosts file. I agree if a bug gets this deep into my system, I'm in trouble anyway. Is this a concern? Oh, that's an interesting way to screw up somebody's system.

Steve: Yeah, it's yet another clever way that Windows could absolutely be messed up. DNS in network configuration can be set to – does not have to be set to your ISP's DNS. In fact, we got a bunch of people on our newsgroups who have found faster DNS servers. And in fact, Leo, you're now becoming a fan of the Open DNS servers.

Leo: I was briefly. Changed away from them, actually.

Steve: Yeah. I have also. So the point is that your computer can be instructed to go to whatever DNS server you tell it to. If something gets in your machine, changes those settings, then, yes, you're in trouble. on the other hand, this is where I remind people, if something is in your machine enough to change your DNS settings, you're already in trouble.

Leo: And if you were on a Mac, it'd ask you for a password anyway before – it doesn't let me change them without asking for a password. That's why security's so important.

Steve: Right.

Leo: Scott Burr of Beaverton, Oregon says: Now that Internet Explorer 7 is out, can we safely go back to using it for browsing the Internet? The reason I ask, I still don't see any new feature that allows us to turn off JavaScript for selected sites. That is, JavaScript can install malware programs, as we've talked about before, through buffer overruns. You're a little more worried about JavaScript than I am. But is IE7 better at protecting against these exploits?

Steve: It is as good as 6, and the mechanism still exists. The way to do this is – it's not clear. And again, it's one of the things I will – it's on my list to get around to documenting, although we have talked about it. IE has this notion of zones. You've got the trusted zone, the Internet zone, the untrusted zone. What you're able to do is, by default all websites are in the Internet zone. That is, they're exterior, they're external. And when you access a website, IE uses the security settings for the Internet zone. So if you turn that security setting all the way high, it will disable JavaScript for all sites on the 'Net.

Now, that's a problem because you may want JavaScript for Amazon or eBay or Google. Certainly you like to have, for example, Google Mail needs JavaScript in order for its Gmail browser-based stuff to function. So then what you do is you place the specific sites that you like and you trust in your Trusted Sites zone, *.google.com, *.microsoft.com, *.ebay.com and so forth. And then automatically, when IE6 or 7 are visiting those sites, the browser looks in the list of trusted sites, sees that you trust this site, and it uses the lower security, the default security, for only the sites you trust. So as with IE6, you're able to do the same thing with IE7. And it is the way I surf. I mean, it is absolutely the way I use IE in a secure fashion. Sites I visit never have scripting enabled at all. And only when I decide to trust them, then I lower my shields and allow scripting to function on the site.

Leo: Well, now we know. In fact, after this Mozilla Firefox thing, I'm thinking maybe I'll

start using IE7. At least it can block third-party cookies without going into some strange config file.

Steve: Yeah.

Leo: Well, that's it for the day. We have 12 questions, 12 answers. I'm sorry to come to the end of it, but it's been a very interesting story. I have to say we thank the folks at Astaro for making this all possible. They're our sponsors on Security Now!, and without them, well, I just don't know what we'd do. Astaro, actually we're just – we love having them on because they're a very exciting company. They do great stuff. They do the Astaro Security Gateway. You've probably heard us talk about the Astaro Security Gateway. If you're a small or medium business network – business, and your network needs superior protection from spam, from viruses, from hackers, complete VPN capabilities, intrusion protection, content filtering, and an industrial strength firewall. Who knew that it could do so much, one little box could do all that. One single, easy-to-use, high-performance appliance, that's the Astaro Security Gateway. You want to try it free in your business, contact Astaro, Astaro.com, or call 877-4AS-TARO. Built on open source, this stuff is really good. And of course if you're a non-business user you could download a free version just for yourself at Astaro.com.

Also want to point you to TWiT.tv/dell. That's where the Leo's Picks live on the special page on the TWiT site. If you're buying a Dell computer, if you're looking for a Dell computer for getting ready for Vista, take a look at the great Dell deals. I'm loving this XPS 700. Oh, man, what a nice computer. But there's lots more there at TWiT.tv/dell.

Steve Gibson, of course, is at GRC.com, home of SpinRite. More happy SpinRite customers to tell us about? Every time we do the show you've got somebody good, saying something great.

Steve: I mentioned – remember last week I mentioned the guy who – for whom it took two and a half days to recover his friend's eight gigs of photos. The drive was completely gone. It was a Maxtor drive that was, like, in really bad trouble. And I mentioned that, you know, even though two and a half days seemed like a long time, you know, compared to sending it off somewhere, it was, you know, substantially less.

Well, I did get a piece of mail. This guy writes: A week ago I started my day with a Blue Screen of Death, advising that I had an unmountable boot volume. Efforts by a Dell tech only lead him to the conclusion that we should reformat the drive and lose all my data. Nothing would recognize the drive, and all of the checkdisk commands in the book could not even see it or result in anything but the same blue screen on every reboot. A Maxtor utility that I ran from a downloaded file advised me to return the drive for a replacement. After getting estimates ranging from $400 to $2,700 to recover my data, and trying numerous other "tricks," in quotes, recommended by online chats, et cetera, I was fortunate to come across SpinRite. At first the glowing testimonials seemed...

Leo: Too good to be true.

Steve: He says: At first the glowing testimonials seemed just too good to be true. And I will admit that I thought they may have even been fake. So I invested the $89, downloaded the file, and fired up the floppy. At first I thought that it was going nowhere because after four hours it still said 2% complete. I figured I would leave it running over the weekend. And imagine my surprise when I came in this morning, saw the message that it had completed,

removed the floppy, and it booted up.

Leo: Wow.

Steve: Ran checkdisk and then started Windows XP. All I can say is wow.

Leo: Wow. That's what I say. Wow.

Steve: He finished, says: Thank you for taking the time to create this program that runs off a low floppy, yet salvaged my 80-gig drive. It's bad enough losing data, but I also saved the hours it would have taken to recreate my desktop, links, et cetera. Needless to say, I'm impressed.

Leo: Cool. Steve, you've done it again.

Steve: Love getting those notes.

Leo: GRC.com, SpinRite, everybody's favorite – my favorite – disk maintenance and recovery utility. SpinRite.info for more testimonials. Of course, when you...

Steve: For real testimonials.

Leo: They're not made up. Not made up.

Steve: Not made up.

Leo: You do also get at GRC.com the 16KB version of the show, the transcription that Elaine does, and lots of show notes. GRC.com. That's Steve's site and the home of Security Now!.

Steve, we've had a great time. I will be back from the cruise next week. We can talk more about security on Thursday. What do you say?

Steve: I think that's a great idea, Leo. We have never missed a week, and as far as I'm – if there's anything I have to say about it, we are never going to.

Leo: I love the...