
# Listener Feedback #83

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

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-230.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-230-lq.mp3

---

**Leo Laporte:** This is Security Now! with Steve Gibson, Episode 230 for January 7, 2010: Your questions, Steve's answers #83.

It's time for Security Now! Episode 230, the first episode of the new year. I won't say the new decade because Steve Gibson's an engineer, and he will quibble. Here he is, Steve Gibson, our host, security expert, the man at GRC.com, the creator of SpinRite. Hey, Steve.

**Steve Gibson:** Yes, are we numbering our years from zero or from one?

**Leo:** We did a decade-ender TWiT last week. And I started off by saying, I know, it's not the end of the decade. But if I do an end-of-decade TWiT in December 2010 I'm going to get far more mail than if I do it in 2009. So I'm doing it now. We moved to two digits, but the decade will end in another year because we started...

**Steve:** Yeah, it's interesting, as an assembly language programmer and actually any programmer, there's always this question of are you counting from zero or are you counting from one? And it's one of the things you learn as you mature as a programmer is that, first of all, that matters.

**Leo:** Yes.

**Steve:** And many bugs have come from someone saying, let's see, if the maximum count is five, then that means I have five things. No, you have six things because zero is

a thing. So...

**Leo:** Programmers learn that pretty early, I think. I know I did. That's like the first thing you learn.

**Steve:** Or they go into the social sciences.

**Leo:** Yes.

**Steve:** They just say, okay, I can't do this. I don't understand.

**Leo:** [Laughing] Steve just burned all you social scientists. Which I am, so there we go. We're going to get to - we have some great questions, I know. You've collated some questions from our listeners, 10 wonderful ones. And we should say, if you have a question for Steve, you can go to his website, GRC.com/feedback and ask a question. Steve loves those questions, and every other show...

**Steve:** Oh, they're so good, Leo. I just - we have got great, great listeners who are, you know, taking the opportunity to send stuff back. I love it.

**Leo:** Now, we've recorded this a little bit ahead because as we are talking, Steve, Kiki and I and the entire gang are in Las Vegas for the Consumer Electronics Show.

**Steve:** Yup.

**Leo:** So if you're listening to this on the 7th or the 8th or the 9th or the 10th, don't forget, you can go to live.twit.tv and watch our live coverage. I'm really trying to bring the big Consumer Electronics Show to your living room, or to your den, or wherever you watch it because you don't have to travel. We're going to show you all the interesting cool stuff. We've got all of our team out there - Ryan Shrout from PC Perspective, Ryan Block from gdgt, Wil Harris, Dr. Kiki, Dick DeBartolo, Paul Thurrott. Of course I'm there. Steve, we're going to launch right into it; right?

**Steve:** So we may not have flying cars, but we have all the rest of the technology...

**Leo:** We're getting there.

**Steve:** ...that George Jetson and Leroy or Elroy, Jane and everybody had.

**Leo:** [Singing] Meet George Jetson. Jane his wife. His dog Astro.

**Steve:** Astro, yes.

**Leo:** You know the Ford Flex has this BLIS system, and it will tell you if somebody's in your blind spot.

**Steve:** Nice.

**Leo:** If you get too close to somebody it pre-tensions the brakes, pre-tensions your seatbelt, I mean, it's - they have automatic parking now. I mean, they really - we're getting very close to automated cars, I think. This is - I think it's pretty exciting.

**Steve:** I'm all for it.

**Leo:** Yeah. We will be safer. Let's get to our questions of the day, starting with Ocie Hudson from Ocala, Florida. He's been watching the packets flow. Sounds like a Crosby, Stills and Nash song. Hi, Steve. In Security Now! #223 - that was the SSL renegotiation bug that we talked about - you remarked, as one of our Q&As from last week asked, if I'm behind a router, and I'm using XP, am I not behind two firewalls? It's like yes, so unsolicited packets are not coming in from outside. Now, that was your response. Here's his story:

I have a Linksys WRT54G router between my computer and the ADSL connection. Windows XP is my operating system. It's running Sunbelt Software Personal Firewall (SPF). We've talked about that before. In fact, I think you recommended it. The router's firewall is configured to block everything inbound, and I have configured the SPF firewall to log dropped packets. Now, here's the thing: During any given 24 hours there are anywhere from as few as 50 to several hundred packets that transit the router firewall and are dropped by SPF. Most, but not quite all, are TCP packets sent by a website which Firefox is currently on or which has been recently connected to a port on my computer that is now closed. Not port 80, apparently. SPF drops them because the port is closed. But why the website sends them, who knows? Why they pass the router's firewall is unclear, but I assume it's probably because the packets have my computer's MAC address in the header; right?

Also from time to time an unidentified process attempts to send an outbound TCP packet from a closed port to some other computer on the Internet, not always the same one. The destination computers have only been identified by their IP addresses. None has ever been identified as hosting a website. So I don't understand why you believe that two firewalls prevent any unsolicited packets from being received by the software and operating system running on a computer. It seems to me that just one firewall should be all that's required. But according to my SPF logging of dropped packets - it can be informative, sometimes mysterious. I wonder whether it might allow some to transit that it should have dropped. In other words, should he be concerned by these results?

**Steve:** Well, this is a great question because it addresses sort of the boundary condition, which I have to actually confess is the reality of at least what his router is doing, which is arguably what it was designed to do and probably maybe the best it can. So he talks about during any given 24 hours there are anywhere from as few as 50 to several

hundred packets that are dropped by - that pass through his router to hit his Windows XP machine where his SPF firewall logs them and drops them. So his question is, okay, wait a minute...

**Leo:** How are they getting there?

**Steve:** Well, yeah. What's going on? Well, the problem with NAT is - we've discussed this extensively and done a couple episodes on it - is that the router sees an initial packet, typically a TCP SYN, that's short for sequence packet, going outbound to a remote website, for example, or whatever you're making a TCP connection to, typically websites, maybe a POP server, SMTP server, whatever. So it creates an entry in a table such that when the packet returns, that it's able to say, oh, I remember seeing sort of the corresponding packet go out there. So now a reply has come back. So that tells it to which computer behind the NAT on the private LAN to forward that packet.

The problem is that routers have different sorts of logic for when to remove that table entry. For example, in the case of TCP connections, there is an orderly shutdown where the computer will send a FIN, for finish, packet. Then the other end is supposed to send back an acknowledgment of that FIN, and then its own FIN packet. And then the computer that initiated the shutdown performs a final acknowledgment of the other end's FIN. So in the same way that there's this so-called three-way handshake to establish a connection, there's a similar three-way handshake to gracefully shut down the connection.

But in order for that to occur, that path has to remain open during the shutdown. That is, if the second the router saw the FIN come in, if it, like, removed the entry from the table, that would cut off the computer from the inside so that it wasn't able to be satisfied. And if the computer doesn't get an acknowledgment of its FIN packet, saying I'm trying to finish this connection, it'll keep sending those, which may recreate a table entry since that's an outgoing packet.

So the truth of the way a router works is a little less elegant than the theory. And so what's happening with the WRT54G is it's deliberately leaving these tables, these NAT-mapping router table entries in place to make sure that both ends are able to satisfy themselves with saying goodbye to each other, essentially. Now, what can happen is that, for example, the far end might not see an acknowledgment, so it'll send a second finish packet. But the SPF firewall, which is - it can be rigorous. When it sees the finish packet come in, it immediately terminates its permission for this connection. So essentially the stateful firewall is tracking the state of the TCP connection. It is very clear and conscious about when the connection is shut down.

So what happens is, as it's truly terminated, that firewall stops permitting packets in that conversation. Yet the router, which can't be that sure because it's not really - it's not the host of the connection. The computer where the firewall is is hosting the connection. The router is just sort of passing stuff through. It has to be more permissive to make sure that packets get through and are acknowledged. So the firewall will emit sort of spurious logging of packets which, I mean, it's exactly what was described here. He's talking about websites he was just visiting, pages he just pulled up, places where Firefox is. And so there was a dialogue and an exchange, it ended, yet the border router continued to let some, like, remaining debris from sort of just the end results of that connection through; whereas the firewall on the computer said, okay, we're done. And then if anything else comes in, it logs it as unsolicited.

So that's what's happening is there actually is a little kind of a gray area in this. And it's worth mentioning that UDP connections like DNS typically uses, since they are stateless connections, that is, there is no startup and shutdown, router logic is very different for UDP. If a UDP packet goes out, a timer is established, sometimes for a minute, sometimes for five, sometimes for 15. It depends upon the logic in the router. There is no notion with UDP protocol of an end of communication. So outgoing packets create these entries in router tables which, as packets continue, that entry is refreshed. And it's only after there's been no activity for some length of time that the timer expires and the router removes that from its routing table. So I would expect if you were doing UDP things you'd probably see even more debris. That's what's going on.

**Leo:** That makes sense. And he shouldn't worry that, because that's happening, that he might be missing other packets, that the firewall's missing stuff?

**Steve:** Well, the one - that's a very good point. The one thing I was going to - I was sort of chuckling when he said, oh, yeah, in 24 hours between 50 and...

**Leo:** 50 and a hundred.

**Steve:** ...several hundred packets. That's like, baby...

**Leo:** That ain't nothin'.

**Steve:** ...if you look at the outside of the router, you want to see something that is effectively blocking stuff? Monitor the outside of the router.

**Leo:** It's 50 to a hundred a minute.

**Steve:** It's a storm out there. So the router is doing the right thing. It is really blocking unsolicited packets. It's just not blocking some of the tailings of conversations that you are having which technically are solicited. Those were dialogues you initiated inside the network. But again, yeah, on the outside, oh, it's bad.

**Leo:** I like the way you said that. It is, it's a conversation trailing off.

**Steve:** Yeah.

**Leo:** Just like [sighing].

**Steve:** Just like [sighing].

**Leo:** Question 2 comes from Troubled. He is a regular in our chatroom, in fact he's in there right now, from Ontario, Canada. He's worried about DNS port randomizations being negated by NAT. And you will only understand what that means if you've been listening to the show since Episode 1. And I'm ashamed to say I kind of understood what it meant. Dear Steve, But not ashamed, proud. Proud.

**Steve:** There you go.

**Leo:** Dear Steve, I just finished listening to your Security Now! episode for December 22. I heard you speak about weaknesses in routers and how you were concerned by the fact that consumer routers seem to prefer to hijack DNS queries. We were saying, why are they doing DNS? Doesn't make any sense. While I think I understand the problem correctly, I'd like to hear your take on it. The specific problem is related to the fact that the recent DNS issue was fixed by making DNS queries use random outbound ports so as to make a MITM be forced to guess - now I'm getting lost.

**Steve:** Man in the middle.

**Leo:** Man in the middle, okay. I've never seen it abbreviated. Okay. So the specific problem related to the fact that the recent DNS issue - is that the BIND problem that we're talking about, that man-in-the-middle BIND problem?

**Steve:** Yup.

**Leo:** Which was fixed in some cases by making DNS queries use random outbound ports so as to make a man-in-the-middle attack forced to guess the source port for the reply. In effect you're getting an extra 16 bits of security, giving approximately 32 bits minus, say, the 1024 ports usually reserved for root. Now, the big question, what happens when you NAT the DNS on a device, as most or all home routers do, that sequentially maps outbound requests? Seems to me you would lose that extra randomness since the router just changes them to be sequential now. While the original query port may have been in fact random, the new NAT'd port is the part the attacker needs to guess, not the original port; right? Of course there's still some guesswork for man-in-the-middle attacks since you'd have to factor in the occasional port incrementing by standard system usage, which could be harder if it's a busy machine. But that's just a few; right? It's not 16 bits worth. Thanks for the shows. Love 'em all. Troubled. That's a good question. You're going to have to explain what it means. I think I kind of get it.

**Steve:** Yeah, it's a good question. And, in fact, exactly that did happen in many instances shortly after port randomization was adopted by servers.

**Leo:** Really.

**Steve:** But it's not the consumer router, it's big iron NAT routers.

**Leo:** It's your ISP's router.

**Steve:** Exactly. So what happened was BIND was updated to issue its queries from random ports. But in the case that the DNS server, the ISP's DNS server was behind their big iron NAT router, then that randomization was derandomized. It was lost. And suddenly queries were going out into the public space from sequential ports. So that had the negative effect of sort of reversing the security improvement of using query source port randomization.

Now, the reason it's not a concern for the consumer, that is, who has a home NAT router that is very likely doing the same thing, is it's not their query to the ISP that we're concerned with spoofing. It's the ISP's DNS server's query out to get the IP from the Internet. So it is the case that, if a man in the middle were monitoring a customer, an ISP's customer's DNS queries, well, then all bets are off. I mean, a man in the middle wouldn't even need to intercept the traffic. He would see a query go out to the ISP, and they could immediately send back a spoofed IP as an answer. So in fact a DNS spoofing in an open network is a complete vulnerability. There's no protection for that.

**Leo:** Okay.

**Steve:** And, in fact, it's funny, as I was thinking about this, I realized that that's obvious, that's another obvious attack, an even simpler attack in any open WiFi network. We've gotten ourselves all worked up about ARP spoofing and knitting ourselves into intercepting traffic and pretending to be the gateway so that we get all the traffic going by. Well, it's like, wait a minute. If there were, like, an evil Google site, then it would be trivial to see anybody sending a DNS query for Google.com out into the, literally into the air, and simply respond with a bad IP. And then that user's computer would then march off in the wrong direction.

So spoofing DNS in an insecure local network is trivial. But that's not where the vulnerability comes from, where the query source port randomness is something that's a problem. It's not the user's query that's being spoofed. It's the user's query goes to the ISP. And that's also, by the way, typically not over the public Internet. It's within the ISP's network. You're behind, if you're on a cable modem or DSL, you have a connection to your ISP. So none of that traffic to their DNS server is visible to the public. It's the DNS server's query outbound that is in danger of being spoofed. And sure enough, a NAT router that's interposed there can derandomize queries. So Troubled is sort of right to be troubled, but sort of in a different place.

**Leo:** Don't be troubled, Troubled. All right. So all right. Now we go to Question 3 from Peter Sinclair. He's in Castle Hill, New South Wales, Australia. And he says: Dear Steve, I've been a Security Now! regular since the start and a SpinRite owner, as well. I heard your recent warnings about overheating in portable disk drives and recalled a paper I wrote - so this guy knows what he's talking about, too - some years ago about disk drive heating effects. It was based on some IBM-sponsored research back in 1989. It turns out - get this. Here's a little math. Heat generation within a disk drive is proportional to the cube of the rotational velocity and inversely

to the fifth power of the diameter. I'd attach a copy but I can't see how on this form. Keep up the good work, and regards to Leo, as well. Peter Sinclair, B App Sc, M App Sc - Bachelor's in Applied Science and Master's in Applied Science. Obviously an expert on this stuff. So the cube of the rotational velocity means it goes up very rapidly with the rotational velocity; right?

**Steve:** Ah, yeah. I wrote back, and I said, "Peter, I have got to see this paper."

**Leo:** Did you get it?

**Steve:** He replied, and I received it, but it was just this morning when I was preparing all of this. So I have not had a chance to review it. And lord knows what the math looks like in there. So it might have taken me a while to digest it in any event.

**Leo:** I love it.

**Steve:** I will make a mention of it next week because the idea that heat generation is proportional to the cube of the rotational velocity and inversely to the fifth power of the diameter, I mean, if I understand that right, it means that, as the diameter gets larger, to the fifth power, heat goes down. Which, okay, I don't know why that would be. I mean, anyway, none of that seems intuitive to me, so I'm really curious. And of course what I do know is that seeking is a big deal, too, because one of the ways that SpinRite tends to heat up drives is that it is, once it starts, it's in there going tick tick tick tick tick tick tick, I mean, it's going step by step by step. Every one of those little cylinder jumps is a burst of high energy because the head is accelerated very quickly and then immediately decelerated. So there's extra power being drawn and extra heat being generated by the drive, which is why SpinRite's also continuously monitoring the temperature of the drive and letting people know, oops, this thing's getting hot. I mean, it will work the drive harder than, I mean, sure, doing a big file copy is sort of the same because the head's jumping around all over the place, writing to the disk. But SpinRite will continue that for hours. So I can't wait to see what the paper says, and I will share it with our listeners. I just got a big kick out that. It's like, here's the equation for heat generation.

**Leo:** I love it.

**Steve:** It's like, whoa, okay, cool.

**Leo:** Here's a question from, or maybe a statement, from Daryl in Kansas. He declares router DNS adds value. We were talking, again, about why these dumb little routers should bother with DNS. Hi, Steve and Leo. Just finished listening to Episode 228 regarding why someone would use a router as the LAN DNS server. I smiled when you brought the subject up. What you say is not useful actually can add a lot of value. In fact, I've locked down my router so that clients attaching to my LAN only use OpenDNS. Now, I know you've heard of OpenDNS. In fact, we talk about it all

the time, OpenDNS.com. There's a big incentive for businesses and families to use OpenDNS to manage network traffic. In fact I do use it at home and here at TWiT.

He says: Incidentally, when I ran your DNS benchmark tool - thank you for a great tool - OpenDNS fared very well. My experience, too. I think it came in second. I really haven't had any problems running this way, and the benefits, as I said, add much value. I can manage my network from any Internet connection and see stats of what's been happening on it. And since I'm a family man, it's been a great tool to filter out the sites I don't want on my network. This guy could be quoting me, Daryl. Anyway, I thought I'd share at least one great reason to use your router as a DNS server. Thanks for a great podcast. I've listened to them all. And the best to you both. Well, that's not what we were talking about, of course, but...

**Steve:** There were two people brought up good points. We've got one a little bit further, I think it's in fact the next one. But I wanted to reiterate this because Daryl's completely correct, certainly. If you configure the router to use specific servers, rather than allowing it to just get whatever DNS servers the ISP serves to it through its DHCP query when it's logging onto the ISP's network, then all of the machines within the network get the use of that DNS just as a function of them doing their own DHCP query during the autoconfiguration of their network connection when they're booting up. You're right, Leo. You had said exactly the same thing. I just wanted to make sure that, since I was so negative on this, that we sort of were on the record...

**Leo:** I'm must be misunderstanding because this is not what the routers are doing, if they're doing DNS. This is just replacing the setting, just as you would on your computer, that's normally provided by the ISP, but explicitly filling in OpenDNS's DNS server. But the DNS is not done by the router, it's done by OpenDNS.

**Steve:** Correct, correct.

**Leo:** When we were talking about...

**Steve:** And that's always the case. Unless you're running your own BIND DNS server which is going out and resolving DNS for you, then the router is just passing this thing through. So you're right, all it's doing is giving you...

**Leo:** An explicit DNS server instead of the one that would be provided by DHCP otherwise.

**Steve:** Right. And it's worth mentioning that this is different than what our next question addresses and what I was talking about because I was talking about the router passing its IP to everyone behind the LAN as their DNS server.

**Leo:** Saying I'm the DNS server.

**Steve:** Exactly, instead of passing the values you configure, that is to say, OpenDNS, onto all the machines behind the LAN.

**Leo:** Right.

**Steve:** Right.

**Leo:** That was what I was understanding is that somehow the router, when you were talking about it last time, actually two times ago, that somehow the router was trying to preempt DNS.

**Steve:** Well, and so the question is, I mean, if you think about it, there are two aspects of this. There's what DNS is configured for the machines on the LAN behind the router, and what DNS the router itself is using. Is it accepting DNS servers from the ISP, which are almost certainly not going to be OpenDNS? Or have you or someone like Daryl gone in and said, I want my whole network to use OpenDNS because it's better DNS? It gives me all these extra features, filtering and monitoring and web safety and so forth.

**Leo:** Okay. Question 5, Tom Zerucha, he's writing from San Diego but living in Michigan. That's just what he says. He has another reason for routers to proxy DNS. Steve, the router has to provide your computer with something in the DHCP field for DNS when you plug it in. And in some situations this could be before it has received its own main address by doing the DHCP query out to the ISP. For instance, maybe you have to configure it with a password or something else. So what does it do? Well, it can't put in an entry for DNS it doesn't have, so it simply implements a forwarding proxy. This makes sense. Now I'm understanding.

**Steve:** Yes.

**Leo:** This is what's going on. This makes sense. It implements a forwarding proxy. When the WAN is configured and/or comes up and the router then gets the real DNS from the DHCP server out there in the world, it can use it. But if the router doesn't yet know what to put into the DNS field for the DHCP request, it needs to send something in that field to the local computer saying give me a lease. I mean, what should it do? How would you solve this? Build a router that is plug and play that doesn't require anyone to manually configure any DNS by telling the computers, just use me as DNS, I'll give you the new DNS when I find out. And then they proxy that off. That makes sense. So is that what's happening?

**Steve:** Well, so yes. So let's back up a little bit. I've used in the last two questions this DHCP. That's the protocol that is used for autoconfiguration of network connections. It's a widely used, universally understood protocol. It's when, for example, in Windows users or on Macs, where there's this obtain an IP address automatically or obtain configuration from the network, what happens is any machine on a network can make a broadcast. So it powers up. It knows nothing, I mean, nothing about the network. No gateway IP. It doesn't know what network it's on. it doesn't know if it's a 192.168 or if it's a 10. or a 172., I mean, nothing. What it's able to do is simply send a broadcast that says, "Hi

there. Is there a DHCP server available that wants to tell me some stuff?"

And so a router, a home network router is, among many other things, a DHCP server. So the computer on the LAN turns on. The DHCP server hears that cry for help and says, yeah. Here I am. What do you need? In the process, it gives its IP to the computer. The computer then says, oh, here's all the stuff I don't know. And basically it's sort of, it's very closely related to DNS, but it can provide other information than just look up IP addresses. Essentially it says, I don't know anything. Tell me what I need to be a citizen on the LAN. And so it receives the IP of the gateway, the DNS servers that it's been told to be configured for. It even gets its own IP. The DHCP server assigns them sequentially or, in some cases, since all of this traffic on a LAN will have a MAC address identifying the adapter, which is to say the computer on the LAN, it's possible to configure the router to always give a certain machine a certain IP, rather than just sort of having them randomly assigned.

And in fact one of the things that machines can do is say, I'd like an IP, and last time I had this one. And so that helps the IP addresses to be relatively static. That is, you're not guaranteed that the DHCP server will give you the same one. But if all the citizens on the LAN suggest what they had before, then generally as they come and go, are powered up and powered down, they'll be reissued the same IP.

Now, this exact same thing happens with the router and the ISP. So when the router is powered up and sees that it's got a connection on its WAN side, not on the internal, on the LAN side, it's the DHCP server facing the ISP, it's a DHCP client. So it does the same thing to the ISP's network that we do to it on our LAN. That is, it sends out a broadcast: "Hi, there. I seem to be on. Fill me up with the specifics of how I should be configured." The ISP maintains a big iron DHCP server, hears the query, and provides all of the connection information to the router.

So what Tom has said, and the point he makes I think is a very good one, is that there are situations like with dialup ADSL, where you've got your LAN connected to the router, which may not yet be connected to the ISP. I'm sort of old school, and I understand how these things work. So I'm always plugging my router into a cable modem or to the DSL router or whatever and sort of letting it get going and stabilize and all that because I understand that it needs to get all of its connection stuff from the ISP before it knows what to give to the network.

Well, clearly, as these connections to the ISPs have evolved and become less static, where you do have DSL connections which are being established and broken, exactly as Tom says, it could be the case that there are computers on the LAN that connect to the router before it connects to the ISP. They ask for all their numbers to be filled in, including DNS. The router doesn't have them. It's got its own gateway address, so it knows to tell everybody on the LAN, hey, I'm your gateway. But it can't pass through any DNS addresses because it doesn't know them yet. So what it does is it proxies for the ISP's DNS, just saying, don't worry, everybody, just use me.

**Leo:** I'm in control here.

**Steve:** I'll take care of your DNS queries as soon as I figure out who I should ask. And so...

**Leo:** Right. That makes sense now, yeah.

**Steve:** It absolutely does make sense.

**Leo:** Yeah. It's almost a necessity that they had to do that.

**Steve:** Yup. It does make sense. So thank you, Tom. That explains why this is being done now. And it would be interesting to see whether it's done as a function of the order in which these things turn on. For example, if the router does have an established set of addresses, and you then connect to it, does it pass the DNS through, or does it always proxy, whether it knows DNS or not? It could be that it's smart, and it's dynamic. That's not something I even looked at before. But it's a great point.

**Leo:** Question 6, Jeffrey Hilgers with the U.S. Navy, Bagram Air Force Base, Afghanistan. Wow, Jeffrey, we salute you in your service to the country.

**Steve:** Yup.

**Leo:** Thank you. He's got a great tip about Microsoft's Patch Tuesday: Steve, in Episode 226, you mentioned in the beginning about where people can go to read about what was released on each Patch Tuesday from Microsoft. I wanted to pass along that eEye Digital Security - which is, by the way, a great, I think they do great work.

**Steve:** Yes.

**Leo:** ...the day after Patch Tuesday releases a very nice bulletin summary of what was released the day before. They also give information on the vulnerability itself and links to Microsoft's information on it. You can view these bulletins at eEye.com/Resources/Security-Center/Patch-Tuesday.aspx. Or just click the - actually the easiest thing to do probably is to go to the eEye website. In the menu there you'll see the Resources menu item. You'll see Security Center under that, and then Patch Tuesday under Security Center. So it's right off the front page, eEye.com, the Resources menu item, Security Center, and Patch Tuesday. That's really a great feature of eEye. I'm glad that they do that.

**Steve:** Yeah. I looked at the page. And, you know, it's a little self-serving. They're saying, of course, oh, well...

**Leo:** We knew about this years ago.

**Steve:** And all of our users who are using such and such, our product, they've always been protected from this preemptively because our Blink technology blinked out before.

**Leo:** I never asked you what you thought about Blink. It's free; right?

**Steve:** It is.

**Leo:** Like personal, yeah.

**Steve:** Yup. And Marc ended up leaving eEye. He's no longer with eEye. He wandered off to pursue...

**Leo:** I didn't know that. Oh, okay.

**Steve:** ...other things. And we've corresponded a little bit since. I think that there are people using it who are very happy with it. There are other people who sort of still feel more comfortable with a traditional pattern-based intrusion detection approach rather than the behavior-based approach that Blink uses. So it's sort of another goody for the arsenal. I wanted - I mention this because I've been aware of this page but never talked about it before, and it is nice. They do a little more open-the-kimono sort of job of telling more detail and more interesting stuff. Microsoft, you could argue, has a little bit of a bias toward minimizing these problems, like I read that one issue about the RPC vulnerability where they said, oh, this could...

**Leo:** Not gonna happen.

**Steve:** This could turn into a worm. And it's called Conficker. And it's the curse of the Internet.

**Leo:** Well, maybe.

**Steve:** Yeah.

**Leo:** Let's not be hasty here. Yeah, it's always true, everybody has their own axe to grind. And EI has kind of the opposite side of it, which is, you know, it's dangerous out there, and use our tool. And Microsoft's...

**Steve:** Yes, we could have protected you from this.

**Leo:** Microsoft's position is it's not dangerous out there. Don't worry. And somewhere in the middle is the truth. But you probably should read both; right?

**Steve:** And even Microsoft's own researchers are saying, oh, don't worry about any of this stuff. Remember...

**Leo:** Wrong. Wrong.

**Steve:** ...from last week the economics of ignoring security advice.

**Leo:** Tim Wells in Marietta, Ohio addresses Firefox update issues. Steve, you mentioned in Episode 228 a couple of weeks back that Elaine, your transcriptionist, had trouble upgrading to 3.5.6. By the way, I think we're at 3.5.7 as we record this. Could be 3.5.8 by the time you hear this. I, too, encountered issues upgrading to 3.5.6 of Firefox on my Netbook. Here's the story. I needed to go to a website, so I grabbed my Netbook, clicked on Firefox. Firefox pops up, says it's updating. So I impatiently wait for my Netbook to update. Then it crashes. And then I couldn't get into Firefox at all. Every time I clicked the Firefox icon I'd get "Mozilla Firefox has encountered a problem and needs to close." So I went to my desktop computer, downloaded Firefox 3.5.6 directly from Mozilla - the full thing, not the patch - and manually installed it on the Netbook, and this fixed the problem. Thanks for the wonderful podcast. Heard every episode.

**Steve:** So I wanted just to pass this on for Elaine's sake. I didn't ever follow up with her to see if she'd fixed the problem. But it's been my experience, within a sufficiently large audience, what happens to one person will happen invariably to many, many more. So if any of our other listeners - it sounds like there was just a little glitch of some sort in the auto update edition of Firefox, but that downloading the whole thing and installing it over what you had works fine and fixes the problem. So Elaine, if you're listening - well, I know you are because you're typing this right now...

**Leo:** Hi, Elaine.

**Steve:** ...and you've just had to type - you've had to type all of this. And this. And this, too. And even this.

**Leo:** And that, too.

**Steve:** Yeah.

**Leo:** That's mean.

**Steve:** And everybody else that this happened to.

**Leo:** Anytime - and I said this two weeks ago. Anytime you install something, a certain percentage of people are going to have problems just from the install because it goes bad. Cosmic ray hits it or whatever, and it goes bad.

**Steve:** It's impossible. I don't know how Microsoft does as good a job as they do.

**Leo:** I believe it. I know.

**Steve:** Or Mozilla and everybody else, for that matter.

**Leo:** Yeah, it's really remarkable, given the hundreds of millions of people who are installing this stuff.

**Steve:** Yeah.

**Leo:** Zec in Colorado is next. He asks about PayPal investigating small purchases as money laundering. Oh, boy. I sell a fair amount on Etsy, which is a handmade online marketplace. To which I give two thumbs up, by the way. Etsy's great. Payment method of choice, PayPal. Today I had a $9 transaction held pending an investigation by PayPal, even though everything appears to be legit, and I even communicated with the buyer. I couldn't believe a $9 transaction would set off PayPal's sirens. So I asked around. One friend said their friend had a bunch of small deposits from writing work withheld, investigated, and ultimately not returned to her because PayPal thought she was laundering money. So what's up? All I do are small transactions with my business. Are they going to keep doing this to me? Keep up the good work.

**Steve:** So this sort of hit me at an interesting time. I'm not sure why I did, but during the Christmas holidays, while I was at Starbucks for a couple hours, from 5:00 a.m. until my family woke up and I rejoined them for breakfast, I happened to read a lengthy, really heartfelt, sad story from a software developer who had really been raked over the coals by PayPal. And we've never really talked about this side of it. I'm an avid PayPal user, as an eBayer who uses PayPal to get money to eBay merchants. And I love using PayPal when it's available because it does prevent having to disclose my credit card information to any sites that offer PayPal. So from an individual user standpoint it's, well, it's a great thing, and I've never had a problem.

Apparently, though, they are so popular that PayPal is a huge target for fraud. And who does that surprise? None of our listeners, I'm sure. The problem is that there's all kinds of automated fraud protection which appears to be hair triggered. Now, I don't know in Zec's case what it was that set PayPal off. In the case of the blog that I read, a software publisher wanted to do a limited-time bundle with some other publishers. And so they agreed over email they were going to be in charge of hosting it. They were a PayPal merchant. Everything was going along. The offer worked really well, which pushed their revenues out of the normal profile that had been established with PayPal and set off warning bells, locked the money, prevented people from purchasing with them. And then PayPal is apparently extremely difficult to deal with.

Again, it's not a business I'm in or I want to be in. So I don't want to be overly critical of PayPal. I think they're offering a service. We've talked about we'd like to see them have better competition so that they do a better job. For what it's worth, the Internet is full of horror stories of merchants, small merchants really being unhappy with their interaction with PayPal. I know that there are plenty of merchants that have never had a problem. Sometimes I think about, for myself, gee, wouldn't it be nice to offer our customers a PayPal button. But then I remember, gee, it's not robust. It's not something I could count on. If I did a new version of - like a new product release, and there was going to be an

upgrade, there would be a flood of retail sales, and it would set off all kinds of alarm bells. And what I've seen, from what I've read, is merchants have no leverage at all. Basically you're just out of luck.

**Leo:** I transfer, any time we get a bulk of money in there, actually PayPal is now my salary. So anytime we get a bunch of money in there we transfer it out because we don't want it to ever get locked up.

**Steve:** Yes.

**Leo:** So we're pretty assiduous about getting the money out of there, getting the money out of there, getting the money out of there. But PayPal, you know, we changed the way the donations, or the contributions, I should say, because it's not a nonprofit work, that's how - that's my entire salary. So whatever people donate goes to me. All the advertising money goes to the company, to hosts, company salaries and all that stuff. I'm the only one who's paid out of PayPal. It's working pretty well. People have been very generous. On to Question 9 from Poojan Wagh in Chicago, Illinois. He wonders, should I be paying for more than 128-bit encryption? He says there's a website, we'll name names, CrashPlan.com. It allows one to do offsite backups to a friend, peer to peer. That's an interesting idea.

**Steve:** Yeah.

**Leo:** I'd like to do this with my in-laws, but I'm wondering if their free offering of 128-bit Blowfish is enough. I trust my in-laws to safeguard the data at their end, but I'm wondering about the data in transit. For a $60 license you can get 448-bit Blowfish. Is it worth it, or is 128-bit enough?

**Steve:** It's a great question. And the answer is 128 bits is more than enough.

**Leo:** More than enough, yeah.

**Steve:** We're talking symmetric encryption. So it's state of the art. Blowfish was designed, as we know, by...

**Leo:** Bruce Schneier.

**Steve:** ...Bruce Schneier, exactly. And it's a cipher which has withstood the test of time - very conservatively designed, well analyzed. It's a very good choice. And this is clearly these guys saying, well, we want to hook people on this...

**Leo:** Free stuff.

**Steve:** ...free stuff. We want to somehow add value beyond what's free. So we're going to give them 448-bit Blowfish. Well, okay. You absolutely do not need it.

**Leo:** This is cool, though, because they can give it to you free because they're not providing the storage. It just - it's like Hamachi or something, connected to another server.

**Steve:** Exactly. They're providing transit. And in fact I'm - now, what's not clear is whether the 128-bit encryption is one-time key or a persistent key. So the reason that Poojan's question sort of says, hey, I'm only concerned about this in transit, is I'm assuming the key is generated with some sort of key handshake, pseudorandomly, and used in a session. But probably another key is used next time you connect. That's the only sane way to do it, in which case 128 bits is just fine. Now, you could also argue that 128 bits is fine forever, which is probably the case. So again, there's just no reason to pay extra money for extra bits of key length on a really good symmetric cipher like Blowfish. 128 bits is fine.

**Leo:** I'm going to have to take a look at this. That's an interesting solution. Hmm. Our last question comes from Matt Ridley. He's in Kaukauna, Wisconsin. I think I just pronounced that the Hawaiian way. Probably not how they pronounce it in Wisconsin. Maybe Kaukauna. He wants some info about - or adds some info about the UAV Predators: Hey, I just listened to the last episode this morning on my way to work. Want to let you know more details about the Predator/UAV news. Two of the guys in an online community that we have shared for eight years or so fly UAVs - that's, what is it, Unmanned...

**Steve:** Aerial Vehicle.

**Leo:** ...Aerial Vehicle, and have at least had some experience with the Predator. Those drones we use heavily in Afghanistan and Iran.

**Steve:** Right.

**Leo:** According to them, after this came out, the following is the current setup: Controls are encrypted, as you mentioned; mission mode cameras are encrypted during flight. This was the issue was can people see the camera output. The unencrypted videos in question are usually from takeoff, landing, and refueling. The reason being, according to them, the pilots, is that the video timing lag caused from encryption/decryption - ah - gives them so much latency that it's not safe during takeoff, landing, and refueling. Well, duh, of course, that makes sense. Once the UAVs are aloft, real-time video feedback is less critical, so encryption is engaged. According to them, in 2012 this video will become encrypted, as well. That makes perfect sense.

**Steve:** It really does. Now, there are many - again, we don't - no one that I've run across sounds to me like they're an absolute authority. So there's lots of people saying...

**Leo:** I'm sure the people who know don't talk.

**Steve:** Exactly. So, I mean, but this explanation makes sense, except that encryption and decryption should be really fast. Symmetric encryption/decryption doesn't take any time unless we're dealing with hardware that is so old, which really could be the case, you know, we run across that from time to time, too, that we're seeing really old hardware which is just - it can't do any more than it was designed to. So I really, from a theoretical standpoint, though, I love this idea that they do not encrypt takeoffs, landings, and refuelings where they really need, for their own piloting controls, real-time video feedback. But once they're in the air, that's less critical, so they click on encryption, get some lag, but have the security that they would hope these systems would have. So that was cool. Thank you, Matt, for sharing that.

**Leo:** That kind of makes sense because even if it's a millisecond of lag, maybe it's a very minor lag, on a critical thing you really want to feel like the controls are real-time; right?

**Steve:** Oh, I mean, any videogame player knows...

**Leo:** Yeah.

**Steve:** ...that if you've got lag between your controls and what you see, it's like, oh, this is just awful, compared to it being really instantaneous.

**Leo:** Well, there you go, 12, or I'm sorry, 10 questions good and true from our listeners. If you have a question for Steve, go to GRC.com/feedback and leave that question for him. I'm sure he'd love to hear from you.

**Steve:** Absolutely.

**Leo:** We answer questions every other show. GRC, when you're there, by the way, is a great resource. 16KB versions of this show, Steve makes those available for easy download. We don't. He does, which I really appreciate your doing that, Steve. He also pays for transcription, another great benefit to the community, and gives away a ton of great free software like Wizmo, DCOMbobulator, Shoot The Messenger, that DNS Benchmark Tool, more and more all the time. Perfect Paper Passwords. GRC, it stands for Gibson Research Corporation, GRC.com.

There's one thing there that you pay for, and I highly recommend you do, not just to support Steve, but because you need it. It's SpinRite, which is absolutely without question the one, the only, the best hard drive maintenance and recovery utility. If you're not using SpinRite, you're not getting the most out of your hard drives. We SpinRite every drive before we put it into service. It's a must-have. GRC.com. Steve, thank you. Welcome to the new decade.

**Steve:** Always a pleasure. As our listeners are hearing this, you're playing in Las Vegas at the Consumer Electronics Show.

**Leo:** Oh, thanks for reminding me. Please go to live.twit.tv for our live coverage and TWiT.tv for the downloads of - we're going to give you wall-to-wall coverage of the Consumer Electronics Show, yeah. And we'll be back next week in-studio to talk again about security.

**Steve:** Episode 231 next week. I'll talk to you then, Leo.

**Leo:** Thank you, Steve.

**Steve:** Thanks.