How Local Area Networks Work, Part 1

Description: Having covered the operation of the Internet's WAN (Wide Area Network) technology over the last two weeks, this week Leo and Steve turn to discussing the way Local Area Networks (LANs) operate and how they interface with the Internet WAN. They address the configuration of subnet masks, default gateways, and DHCP to explain how packets are routed among machines and gateways within a LAN.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-027.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-027-lq.mp3

Leo Laporte: This is Security Now! with Steve Gibson for Thursday, February 16, 2006, Episode 27: How LANs Work. Steve Gibson, hello. Leo Laporte here. Steve's in beautiful Irvine, California, via actually Google Talk because Skype wasn't working for us for some reason.

Steve Gibson: Yeah. Let's hope this comes out okay.

Leo: So this is a couple of days after Patch Tuesday, Microsoft. I know because when I turned on my PC it said, "You have updates."

Steve: Right. Yeah, Microsoft released seven patches. Two they marked as "critical," and five were just "important." We know, of course, from our coverage of the Windows MetaFile, where we really looked carefully at what these designations mean, we know that "critical" technically has to be worm-able, you know, an exploit that really could damage the Internet on the scale of a worm getting loose. And anything that isn't that, they're calling "important." So there was actually some follow-on Windows MetaFile stuff, and something to deal with the Media Player able to be tricked into showing an image or something. So anyway, they're, you know, important. And anybody who's keeping their Windows machines up-to-date will definitely want to take a look at that.

Leo: Now, here's the question. When you have Automatic Updates turned on, will it only download the critical, or will it download the important, as well?

Steve: Oh, it gets them all.

Leo: Okay. It does say "Check for important updates and install them." So, and that's what you'd want to do, it sounds like, since they've kind of changed what "critical" means, is you want to get everything that's important.

Steve: Oh, yeah. I mean, it blew my mind that they weren't considering the Windows MetaFile problem an important problem - or a critical problem. They were only calling it important. So you definitely want all of the updates.

Leo: Given that naming, we definitely want to get the important files, as well.

Steve: Well, so you're back from your cruise.

Leo: Yeah, and thank you for "Dragon Fall." What a wonderful...

Steve: "Fallen Dragon."

Leo: "Fallen Dragon," sorry. Can never get that name right. That's the Peter Hamilton book that you sent me. And I am loving it.

Steve: Good.

Leo: He's - great sci-fi. Good, hard sci-fi. Had just read Stephen Baxter's "Ring," which is very hard sci-fi. It's all about string theory. So this is a little breath of fresh air because it's written - the characters are really deep and rich, and the worlds are very rich. He's a great writer.

Steve: I love his work, yeah.

Leo: And you just saw "Firewall." Now, let me ask you - this is the new Harrison Ford movie. Does it have any firewalls in it?

Steve: No.

Leo: Oh, rats.

Steve: I don't know why, I guess they used the word just because it's kind of a cool word.

Leo: It's hot right now.

Steve: Maybe, you know, it is about banking security and computer networks. They've got sort of some quasi-believable technology in there, you know, it's hokey, you know, borrows his daughter's iPod to download some account numbers and things. Like, oh, okay. But Leo, it's a fantastic movie. I loved it. I guess the critics didn't like it that much. But, I mean, I was on the edge of my seat. So, I don't know. I recommend it. It was really good.

Leo: And no big technical mistakes in there, huh?

Steve: Ah, they sort of avoided all the technology stuff. There was some stuff that, I mean, you know, you just sort of allow it in a movie like that.

Leo: Yeah, yeah.

Steve: If anyone wants to get their technology, I think they have to keep listening to Security Now!. They're not going to get it from a Harrison Ford movie.

Leo: Willing suspension of disbelief, I think they call it.

Steve: It was great, though.

Leo: So let's cover any issues from previous podcasts before we get into this week's subject, which is - we've talked about how the Internet works. Now we're going to talk about your network, how your home network works.

Steve: Yeah, we had a couple people sort of took me to task on the issue of UDP versus TCP. You know, I was talking about, essentially, what a nice protocol UDP is. And they were saying, hey, wait a minute, you know, TCP is so useful. And just a little bit of a clarification. You know, like I've mentioned how DNS uses UDP because the responsibility for making sure a reply comes back is the application's. The real way to think of UDP versus TCP is in terms of overhead. For example, a DNS query is one packet that goes to the server. And the DNS reply is one packet that comes back. So the entire transaction for making a DNS query over UDP is just two packets. By comparison, if you were to use TCP to make the query, and in fact DNS does support TCP also, you'd have to send one over - well, instead of enumerating them, because we did that last week, you'd have to do the TCP three-way handshake, so that's three packets; then the query and the response, so now we're up to five; then two packets to shut the connection down. So that's seven packets to make a query over TCP which you can do in two packets with DNS.

Leo: With UDP. With DNS? UDP.

Steve: I'm sorry, with UDP, exactly. So probably the best way to characterize it is saying, you know, for short things or for real-time things, UDP makes more sense. For longer, persistent connections where you want the network underlying layers to worry about getting the stuff there, getting the packets in the right sequence, and all the things that TCP does, it definitely makes more sense, and all that work is done for you.

Leo: All right.

Steve: And another interesting note is some people have been reporting that GRC's ShieldsUP! system has been showing them some open ports where they used to be all stealth. Well, in working back and forth with them, we've tracked it down to the growth of UPnP-aware peer-to-peer systems. And this is something that you and I, Leo, I have to say we were ahead of this by, you know, a year or so. I knew that Universal Plug and Play enabled in routers was going to be a problem. And what's happening is the various peer-to-peer systems are now beginning to take advantage of that to open ports through people's routers in order to sort of, you know, solve the interconnection problem. But at the same time it is lowering their security.

So, for example, BearShare and LimeWire actually have verbiage on their sites that says "UPnP support can now detect enabled routers and automatically configure the router for LimeWire's optimal performance." Even the BitSpirit, BitTorrent client, it says "BitSpirit is a powerful and easy-to-use BitTorrent client. It supports simultaneous downloads, download queue, UPnP port mapping, and NAT traversal." So what we're beginning to see is we're seeing these clients are, behind people's backs, without notifying them, reprogramming their NAT routers to open ports back through the router. That's something, I mean, sure, if it's what you want, and you know it's what's going on, then hey, that's cool. But what's happening is this is being done, and people are running ShieldsUP! and saying, wait a minute, why do I have ports open? So this is happening behind people's backs.

Leo: They're being opened by these applications.

Steve: They're being opened by the applications because the routers have Universal Plug and Play enabled in the router.

Leo: And this is what we've been warning people against, and saying turn that off. Of course, it is a convenience, isn't it, when LimeWire automatically turns it on. Then you don't have to even think about it.

Steve: Well, it certainly is a convenience. And again, this is - we can now expect that trojans are not far away from doing this.

Leo: That's the problem. Convenient for LimeWire means convenient for trojans.

Steve: So, you know, in terms of progression from a security standpoint, the feature was created which, you know, we quickly identified as a bad idea. Then the feature got used by good programs. And we can expect shortly to see the feature being used by bad programs.

Leo: Great.

Steve: So again, you know, it's an early heads-up to people.

Leo: That's, yeah, wow. All right. So it's ShieldsUP! saying, yeah, you do have ports open, you just didn't know it. Surprise.

Steve: Well, exactly. The other thing I need to talk about, I do want to mention briefly that Mac OS X has had the appearance of its first worm written for it. It's an iChat propagating worm. It goes and enumerates all the applications that the user has used in the last month and replaces them with copies of itself that then links to those applications. So, you know, we are beginning to see, as the Mac popularity increases and as, you know, as it spreads, we're beginning to see, you know, the growth of security issues over on the Mac side, too.

Leo: Yeah. Although this is, as you said, it's not a very serious attempt at an exploit. In fact, I don't think it's available in the wild. It's supposed to spread itself through iChat, but I don't think anybody has said it has. And in order to get infected by it, you would have to unzip the file, open the application, which claimed to be - so if you're not paying attention, I guess, but it claimed to be pictures and is an application, okay. You open it. And then it asks for your administrator password, which you would have to give it.

Steve: Right.

Leo: So you have to really cooperate if you're going to get infected by this thing.

Steve: And lastly, people are wondering where we stand with our promised OpenVPN How To guides. That's all I'm working on right now. I've got OpenVPN servers running at our facility at Level 3. I've got one here. I just converted one of those, the local FreeBSD box, to diskless operation, so it's running with no moving parts, no fan, it boots off Compact Flash and works perfectly. And then just yesterday I set up a Windows 2000 machine to work out the final details of OpenVPN on Windows 2000 and XP as the server endpoint. So I want to assure people that, you know, we lost a month, basically, through the whole Windows MetaFile thing in January. But I am back working on the OpenVPN How Tos, and we will do one final show to introduce those and talk about those once those are all up and posted. But they're coming.

Leo: All right. There's still a lot of interest in that. We haven't moved on completely from it.

Steve: Oh, there's a ton of interest, Leo, because OpenVPN is such a beautiful solution for allowing people to access networks, you know, from a road warrior profile, when they're on the road. And it's free.

Leo: Great. So, on to LANs.

Steve: Yeah. We've talked of course now two episodes about the WAN side of this networking technology, about packets and ports and IPs, and then the ICMP, UDP, and TCP protocols. We're going to do two episodes on LAN technology. Next week, of course, we've got our Mod 4 Q&A episode. And then the one after that I want to talk about the ARP side of what glues IP and Ethernet together, and specifically the security concerns for that. But first I want to lay some foundation with sort of like what goes on on your LAN, what is

a default gateway, what's a subnet mask, and how does all that work to essentially handle and route packets around the Local Area Network.

Leo: It's been interesting to see the change that's occurred. It used to be networking was really the province of one guy at the office who knew this, and nobody else did. And now everybody who has a computer at home is learning about subnet masks and DHCP. It's become part of our vocabulary.

Steve: Well, yeah. And of course, you know, the Internet technology itself is just spreading through the lives of everyone who's associated with it. Now, prior to routers happening, that is, the whole personal NAT router, people would normally just connect their computer directly to whatever Internet connection they had. Their modem would dial them into a service provider; or if you had DSL or even a cable modem, your computer would be connected directly to it, and so you'd have less contact with this sort of LAN technology because your machine would be part of a LAN that was extended to you by your provider.

But now, with NAT routers, which we promote so much because they're such great security, essentially the router is creating a network of its own. And so there's more of this LAN glue that's exposed to the user. The whole key to a LAN that's connected, essentially, to another network, to a non-directly connected network, is this router that we've talked about. We saw how Internet routers out on the 'Net are sort of out there in the Internet cloud, moving packets closer to you. Well, the router that most people have direct interaction with now is their own residential NAT router.

And so essentially what happens is you've got one or more machines that are all connected together by a hub or a switch. And by connected together, essentially it means that they're on the same network, that they're able to address their packets directly to each other. The question, though, is how do they know that the packet is on an IP on their own network, or that it needs to go to the router for remote routing? Well, that's this thing that's called a subnet mask. It's what the subnet mask does. And, you know, people who are using normal NAT routers will typically see - the IP of their router will be 192.168.0.1, for example.

Leo: That's the gateway router.

Steve: Well, exactly.

Leo: Okay.

Steve: And that'll also be the IP that their router participates in the Local Area Network from. Now, the subnet mask of that can vary. But, for example, it's often 255.255.255.0. Or it might be 255.255.0.0. But either way, the idea is that the IP that a packet is destined for, that is, the target IP, is compared against this subnet mask to see if the digits where we've got 255s are the same as the machine's own IP, essentially. The idea being, for example, that a computer connected in the LAN might be 192.168.0.2, for example. If it's sending a packet to any IP with the last byte different, so like .2, .3, .4, where the left-hand side of the IP has all the same numbers, 192.168.0, well, it knows that, because of the subnet mask, which is saying, essentially, specifying which IPs are in the local network, it's saying, okay, you know, this machine is directly reachable by me. So it addresses the packet to that machine that it's sending the data to directly, as opposed to saying, for example, if you were connecting to Microsoft, and you had some completely different IP, you know, 4.79.something else, it would notice that the digits were different where the subnet mask was a 255.255. That's sort of the whole key to this. That tells it that this is an IP not on the Local Area Network, but one that is located somewhere else.

Leo: So 255, if you put it in binary, it's maybe a little easier to understand. Each of these is eight bits. And a 255 is all ones.

Steve: Correct.

Leo: And a zero is all zeroes - eight bits of ones, eight bits of zeroes. So it would be 11111.11111.11111.0, or .00000. So where it's ones, that's where you pay attention. Where it's zeroes, it could be changed.

Steve: Exactly. So I'm glad you said it that way because it's - where it's ones, you compare your destination IP to your own network to see whether they match up. If they differ, then that says, oh, this is not an IP on our Local Area Network. We need to send this to the gateway.

Leo: Maybe another way to understand it is that allowed addresses on the LAN, on the local network, are those that are not masked by ones.

Steve: That's another way of saying it, yes.

Leo: So if it's got zeroes in those places. Then that would mean that it doesn't always have to be 255 or zero; it could be 254 or some other number. Yes?

Steve: Well, that's actually true. And it does get complicated. While I was trying to figure out how to visually describe this, I thought, well, you know, if I get into networks where the mask is sliding to different values, we're really going to get ourselves tangled...

Leo: It does get complicated.

Steve: Tangled up and lost. But, for example, the 192.168, we've talked about this network before because it's been reserved so that it's only for use in private networks. Another network, for example, is the 10. network that we've talked about, where it's 10.anything.anything.anything. Well, the subnet mask for a 10. network would be 255.0.0.0, meaning that any IP beginning with 10 is going to be part of that same 10. Local Area Network. And then all the IPs beginning with 10 having anything else fall within that. But anything that doesn't start with a 10 will be outside of that LAN.

Leo: So when you put a subnet mask into a router, you're telling the router which places are significant when it comes to routing versus local. Where to pay attention, essentially.

Steve: Well, exactly. And all of the machines on the same LAN will have the same subnet mask.

Leo: They'd have to, yeah.

Steve: Yes, yes. There is an agreement among them all that we're all machines in a network that contains this block of IPs and doesn't contain all the other IPs. And so that's the key. It's like, is this IP in our LAN, or is it not in our LAN? Now, if the IP that a packet is being addressed to is not in the LAN, all of the machines will send their packets to the so-called "gateway machine."

Leo: It must be out there somewhere, so here, you take it.

Steve: Well, exactly. And essentially, what's really interesting - and this is what we're going to cover in two weeks is a little more of the mechanics of that, when we talk about how Ethernet works and what the whole ARP, the Address Resolution Protocol is. But the idea is that the packet is addressed to a remote IP but sent to the gateway. So the gateway, for example, a NAT router in most people's experience, will - it'll receive this packet. So it actually receives it. But when it looks at it, it says, wait a minute, this is addressed to an IP outside of the LAN. So that's why the router has the same subnet mask as all the machines do in the LAN. So it does the same comparison. It sees that this packet is addressed outside of its LAN, and it says, oh, that means I send this out the WAN interface.

Leo: There's another side effect because it will also say how many available IP addresses you can use. For instance, if you have a mask of 255.255.255.0, it means you have a total of 255 addresses you can use in that network.

Steve: Yes. You lose a few. You lose - the very last address in any LAN is the so-called "broadcast address" because, by agreement, a packet addressed to that is actually heard by everyone on the LAN. It's a broadcast. And in fact there have been some early exploits that have taken advantage of that to do all kinds of, you know, wacky things with ping packets. And then the zero address, that is, the zero IP is also reserved as the so-called "network address," that is, the address for the whole network. And then, so like you have - so you'd have IPs 1 through 254. One of those IPs is going to be the gateway. Unless you, I mean, you could have a LAN with no gateway. If it was just a LAN not connected to anything else, you could have a bunch of machines all connected just to a hub or a switch that are not connected to the Internet, in which case they wouldn't have a defined gateway because there's no way for them to send traffic outside of the LAN.

Leo: So you really effectively have 253 addresses.

Steve: Exactly.

Leo: Yeah.

Steve: And now many - there are some routers that will allow you to configure all of those. Other routers will say, well, you know, no one in a residential mode is going to need more than, you know, 100. And so many routers do restrict you to a lower number overall.

Leo: And then they talk about Class A, Class B, Class C networks. That's just how many addresses are available for that network.

Steve: Well, yeah. And in fact, as we talked about two episodes ago, the original concept of this was that networks would either have - would be 255.0.0.0, which is a Class A network, like the whole 10. network is; or a Class B network would be 255.255.0.0 for their subnet mask, in which case you've got 16 bits' worth of machine identifiers down in the right-hand two bytes; or a Class C network would be 255.255.255.0, in which case, as we just said, you've got essentially 255 or 256 possible addresses, but you lose a couple of those to specific defined applications.

Leo: Ah. You know, I've been using subnet masks for so long, and now to know what they really mean, it's so nice.

Steve: Well, it's an important aspect of our networking. Now, many people will just use - they'll configure their Windows or their Mac or their Linux client to so-called obtain an IP address automatically. And essentially what that does is, when the computer is turned on, it will send out a broadcast saying, hi there, I'm new in the neighborhood, literally. Is there a DHCP server listening that will hand me an IP? And so they're able to, knowing nothing about the network, to broadcast this DHCP query, saying has this LAN got someone who can give me an IP address. And in the case of a residential network, that'll be the router that'll be turned on and waiting. It'll receive this broadcast because a LAN - and we'll be talking about this in two weeks - has the ability to essentially have, like, an all-stations-receive sort of deal. So not even - knowing nothing about the network at all, whether it's a 10., 192.168, I mean, or any IP, it's able to say, hey there, can somebody help me?

The DHCP server or service process, which is running in the LAN's router or gateway, it will respond to that and say, hi there, I'm responsible for doling out IP addresses. And so it'll establish a dialogue with the machine that has broadcast the query network-wide. And so it'll say, here's what you want to use for your DNS servers. Here's what you need to set your subnet mask for. And here's my own IP to use as a gateway. And here's an IP you can use that is unique for yourself. So it's the DHCP server in a network that's configured dynamically which is responsible for doling out these IPs as machines first appear on the network.

Leo: So when you set your Windows box to obtain IP address automatically, that's the setting in the network settings, that's what it's doing. You're saying use the DHCP server and get a number from them.

Steve: Right. And as you might imagine, I mean, if you think about it for a second, since that's a broadcast,

and everyone on the network hears it, if somebody malicious were to hear that, they could stand in for the DHCP server and completely fool a machine into essentially trusting it as the gateway instead of the machine.

Leo: Now, a router can still act as a gateway even if it doesn't do DHCP. You could easily set up what they call "static IP addresses" where you give each machine its own address - you have to make sure they don't conflict - and turn off that DHCP functionality. But you'd still have routing; right?

Steve: Yeah, I think that most routers do give you the option of shutting down DHCP. What that then means...

Leo: Would that be more secure?

Steve: It would be, well, I would say yes except that we're going to see in two weeks how there's really no semblance of security on Ethernet networks at all.

Leo: So forget it.

Steve: I mean, it's a disaster, Leo.

Leo: Oh, okay.

Steve: So you wouldn't really gain anything.

Leo: You don't have to worry about somebody spoofing your DHCP router.

Steve: As far as I know, there are other ways to achieve the same things that are even more immediate. Because spoofing your DHCP would require, I mean, there are all kinds of - basically there's no security at all on an Ethernet network. It was...

Leo: And we'll talk about that in a couple of weeks. I know, it's a broadcast network and everything.

Steve: Yeah, it was just - there's just no security at all. So the reason someone might do that, though, is if, for their own purposes, they didn't want the IPs of their machines to float around. For example...

Leo: Or you didn't want somebody to join your network promiscuously. Right? If there's no DHCP server, they - oh, I guess they could. They could just...

Steve: Yeah. All they would have...

Leo: ...pick a number at random, huh.

Steve: Exactly. Or not at random. But if they were able to look at the network traffic, they would see what range of IPs you were using and be able to just give themselves an IP that's not currently in use and immediately be able to participate as a peer on the network. But, for example, say somebody wanted their, like - they had, like, TiVos. I have a pair of TiVos that are on my LAN. And I've given them static IPs. Actually I'm not using DHCP in my own network at all because I know what all the IPs are. And I sort of have, like, my servers are 10.1.something, and my TiVos are 10.0.0.200 and 10.0.0.201. I just, you know, I sort of have my network blocked out in large chunks and arranged in a way that makes sense to me.

Leo: I've run into people working on their networks where it's set up statically. Is there an advantage to doing it? Is there less overhead by turning off the DHCP server? Is there any reason you might want to do that?

Steve: I would probably have to qualify that as sort of the propeller head approach.

Leo: It is, because you have to constantly kind of pay attention. You can't just give, you know, give somebody an address. You have to make sure it's not conflicting with another one.

Steve: If someone comes over to your house and wants to, you know, plug into your network, you know, certainly their machine, if they are normal, has been just left in the default mode of obtain IP address automatically. You know, they do that at their house. They come over, plug their computer into your network, and it doesn't work.

Leo: That's when I got involved with this guy's network, and I was really sorry, too. I said, what's going on? And I figured, oh, you have static IP addresses and blah blah. So is it an advantage? It's not really an advantage to do it that way.

Steve: No, I really see, I mean, there's really nothing wrong with using DHCP. It's no less secure, and I hope I didn't scare people away by talking about, like, you know, spoofing a DHCP server. I didn't intend to. It's certainly possible.

Leo: It's not likely.

Steve: Well, as far as I know, it isn't a common exploit; and there are far more common exploits against Ethernet networks, you know, against the LAN networks we're talking about, than doing that, that are far more effective. And we'll be talking about that in two weeks. So I think DHCP makes a ton of sense. I'm just, you know, I'm someone who likes to know the IP of every machine I've got. And I...

Leo: I bet you - I'm just guessing. Never been to your house. I bet you actually have a little DYMO label on each one with the IP address on it.

Steve: You know, I did, except that they started getting obsolete. And the only thing worse than a machine that is not labeled is a machine that is mislabeled. Because, you know, for some reason you change the IP to something, and now the label is referring to the wrong IP.

Leo: I can just see all your machines with little addresses on every monitor. You knew exactly which one was which.

Steve: Yeah. Pretty much I have them all memorized.

Leo: You know, yeah.

Steve: Yeah.

Leo: Okay. So we've got subnet masks, we've got DHCP.

Steve: We've got the default gateway...

Leo: Gateways, yeah.

Steve: ...where packets are sent. If they don't match up with the subnet mask, they are sent to the gateway machine at the gateway IP. It receives them. It sees the same thing. It sees, wait a minute, this isn't part of the LAN, because otherwise it just wouldn't do anything with it. It wouldn't send it on. But it sees that it doesn't match up with a subnet mask for the digits that are significant, where the subnet is ones, as you put it. And so it sends it out to the WAN. Similarly, when packets come into it, it does the same thing again. It receives the packet and does its NAT translation to translate it into an internal IP and says, oh, look, this does fit within the subnet mask, so it rebroadcasts it out of its LAN side so that the properly addressed machine in the LAN is able to pick up the packet. And that's how all that works.

Leo: Ah. It's really amazing. It's actually, again, and I said this when we described how the Internet works, it's really kind of a remarkable little thing going on behind the scenes. Most people don't pay attention to it, don't even have to. You just plug in your router, and it pretty much just works out of the box.

Steve: It's just elegant technology.

Leo: It is, yeah.

Steve: And of course, you know, in terms of DHCP, since most ISPs use the same DHCP to assign IPs to their residential computers, the NAT router itself is a DHCP client. When it turns on, out on the WAN port, going out toward the Internet, it will say, hey, I need an IP. And so, again, the ISP will be running its huge DHCP server, which is providing IPs to everybody plugged into its whole network. And so it will assign the router a unique IP for the WAN connection that the router has, just in the same way that then the router turns around and is the DHCP server for your own local LAN. So the same thing is going on on both sides. So your router is a DHCP client for your ISP's WAN side, and it's the DHCP server for your little residential LAN side. Very cool.

Leo: Yeah. And I guess you could go on and on and on. Who knows? You might have a - the ISP might have its own DHCP server.

Steve: Yup.

Leo: Well, there we go. I think that was - that's it, huh? There's pretty much the story of your home network, your LAN.

Steve: Yes. It's the first chapter. We're going to talk about the ARP side and the lack of Ethernet security in two weeks, since next week is our Q&A episode.

Leo: So Part 2 on Episode 29, but 28...

Steve: Is what can go wrong.

Leo: What can go wrong with this elegant solution that seemed so good at the time? Steve Gibson, always a pleasure. I learn so much from every one of these. And I know that many schools and universities are using them in their classrooms now. And lots of people like the transcripts, too. You can get those from Steve's website, GRC.com/securitynow.htm. You'll find 16KB versions for the bandwidth impaired; transcripts of every show, thanks to Elaine; and of course Steve's show notes, often with diagrams and further information and links. And that's a good place to go, too, if you ever have hard drive woes, because SpinRite is Steve's day job - although I don't know, he's spending so much time on Security Now!, it might be your late-night job. I don't know when you find time to do it. SpinRite is the

ultimate disk recovery and maintenance utility. I've been using it for years. Version 6 is just out and just fantastic. In fact, our neighbor boy came over, Stephen came over a couple of days ago and said, "My system's acting flaky." And I said, "Son" - I hope you don't mind, Steve...

Steve: No.

Leo: "Take this SpinRite and use it." I felt like such a...

Steve: You have an implicit site license, Leo.

Leo: Well, I'll get it back from him. Don't worry, he won't keep it.

Steve: Ah, good.

Leo: But I'm hoping it will recover his system. I'm sure, if anything can, it will. It's a wonderful program everybody should have.

Steve: Does it for lots of people every day.

Leo: GRC.com. Steve, enjoy the beautiful Southern California weather. While we're frying, they're freezing back East. But that's the way of the world, I'm afraid.

Steve: Yeah, this time of year.

Leo: We'll see you next week, next Thursday. We'll be in Toronto, and we'll be doing our every four weeks Questions & Answers. Do people have a place to go on your site where they can ask questions for that?

Steve: Oh, I am just swamped with questions, yes.

Leo: We don't need more.

Steve: At the bottom of the SecurityNow.htm page is a form that people can submit questions to.

Leo: Great. And we do appreciate them. They help us, well, they help us a lot, guiding the show and the content of the show, as well as...

Steve: It gives me a lot of feedback. You know, we had just great response from these first two episodes, sort of our technology foundation series that we're doing now. And we've got a few more things to cover, and it's looking good.

Leo: Right. Thanks to our friends at AOL Radio for broadcasting this on their podcast channel, and of course for providing us with the bandwidth so that you can listen to Security Now! each and every week, absolutely free, whether it's on iTunes or with iPodder X or iPodder Lemon or whatever client you use. We're glad you listen, and we hope we'll see you next time on Security Now!. Take care, Steve.

Steve: Thanks, Leo.