Internet Proxies

Description: Steve and Leo discuss the entire range of applications for Internet Proxies and Proxy Servers. They describe the many different uses for proxies while discussing both the benefits and the potential security and privacy liabilities created by filtering and caching web and other Internet content.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-062.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-062-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 62 for October 19, 2006: Proxies.

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!, our regular visit to the deep, dark underside of the Internet, Windows, Mac, and Linux. Steve Gibson, our security wizard, is here from GRC.com.

Steve Gibson: Hey, Leo.

Leo: Hi, Steve. Boy, you know, once again we were ahead of the curve last week, talking about these Windows exploits, the 0-day exploits that come out right after the first – or the second Tuesday of the month.

Steve: Yup.

Leo: And there was another one, a PowerPoint exploit, fortunately which doesn't seem to be out in the wild. It's just – it's amazing when these patches come out. There were 23 patches last Tuesday. And boy, they just – the hackers get to work.

Steve: Well, and in fact I'm glad you brought that up because I wanted to reiterate that there is still this unpatched direct animation path control exploit that we talked about last week.

Leo: Right, right.

Steve: It's – Microsoft has not addressed it. It's not clear whether or when they're going to. And there is public exploit code in the wild. I shot a note off to our friends down at Sunbelt Software, asking whether they had seen – whether their web spiders had encountered the active exploitation of this in the same way that it was like a wildfire with that VML, you know...

Leo: Right, right.

Steve: ...the previous 0-day. I've not heard back from them yet. But I do want to remind people that, you know, as we talked about last week, if someone, you know, didn't get around to doing – using the little ZERT utility to temporarily disable the handling of direct animation in Internet Explorer, if they use IE, if they have scripting enabled by default, then you really do – I think it's worth just going and using this little bit of free software to turn that off until Microsoft patches. And we can assume it'll be in the second Tuesday of November that they will be patching. So, but that still leaves, you know, four weeks of open vulnerability window. And, you know, the way things are developing now, the malicious hackers are deliberately releasing new exploits at the same time Microsoft patches to optimize or maximize the amount of vulnerability. And it's all about installing junk in people's machines. You know, junk that no one wants installed on their machines.

Leo: Yeah, yeah, it's just amazing. And the fact that Microsoft isn't responding to this problem, this DirectX problem, surprises me, frankly. I guess they're still, you know, analyzing the issue.

Steve: Ah, yeah. And in fact, another interesting little bit of news is Vista related. I'm sure you've probably picked up on some of this, Leo. The security companies have been very upset with Microsoft because Vista protects modifications of the kernel, which, you know, programs like McAfee and Symantec need to have access to the kernel in order to install their hooks. Well, Microsoft's Vista, one of the new security enhancements is a, you know, explicit protection from making any modifications to the kernel. The problem is, you know, from an antitrust standpoint, this locks out third-party software, which has been able to use, basically, their own benign hacks in order to hack into the XP kernel to install themselves at the low level that they need to in order to monitor what's going on. So what's real – I kind of got a kick out of this, is Microsoft has said, okay, well, um, yes, we see that could be a problem. So how about this? In the Service Pack 1 release we will provide the API, you know, brand new APIs. Microsoft apparently is just weeks away from what's called RTM, Releasing to Manufacturing, the final version 1.0...

Leo: Right, October 25th, it's one week away, yeah.

Steve: Yeah. And so Microsoft is saying, well, there aren't APIs in the kernel now that would allow this. And so we can't put them in in time. Because, I mean, after all, it takes Microsoft a month just to patch a flaw.

Leo: Right.

Steve: So, you know, you could argue that there's no way, you know, they could delay Vista's

release. So it seems unfortunate that, I mean, essentially what's going to happen is, as we know, a major service pack release is typically, you know, 12 to 18 months downstream. So none of these third-party tools can be developed and used until Service Pack 1 of Vista. Which means that basically they're going to shut out the market for third parties. Everyone will end up using Microsoft's new security solutions and tools, which are Vista compatible. And as a consequence of this, it's very likely that, you know, these traditional security companies are going to, you know, not be able to hook in their tools, and probably never end up achieving the market share that they would have if they'd been able to release their stuff right at the release of Vista.

Leo: Well, you know, it's complicated because you can of course say, well, this is Microsoft trying to assert a competitive advantage. But I have to think that, given the trouble they've had with antitrust and the Department of Justice in the past, that they'd be very unlikely to attempt that. And they can make a pretty good argument, well, the reason we don't release, you know, the reason we put this patch guard technology in is because it's not – we're not worried about Symantec and McAfee, we're worried about hackers.

Steve: And Leo...

Leo: And having an API makes it open to hackers.

Steve: And, exactly. I would rather Microsoft did not open it because I think the danger of them doing that, I mean, you know, the first question that pops up is, okay, wait a minute, if you create...

Leo: If there's an API...

Steve: Yes.

Leo: ...anyone can use it.

Steve: Then the bad guys can use it, too, exactly. So, I mean...

Leo: I mean, it's a win-win for Microsoft. It both gives them a competitive advantage, and it makes their system more secure. By the way, I'm just reading the article, it looks like that the API will only be for 64-bit versions of Windows anyway. It won't be...

Steve: Right.

Leo: ...for 32 bits. So...

Steve: Right.

Leo: And that's of course where the big money is, is in the 32-bit version.

Steve: Yeah. So I think we're going to see very slow adoption of 64-bit, just because, you know, the hardware is not there. Although you could argue at the same time that Vista is so resource-hungry that we're going to need a lot of...

Leo: All that RAM, yeah.

Steve: ...I mean, you know, basically new machines anyway.

Leo: Yeah.

Steve: Although the 64-bit platforms are probably going to be a lot more expensive than the, you know, the much more mature, you know, faster, you know, Pentium 4 machines.

Leo: Now, today on the show we're going to talk about proxies. And I want to do that in a second. But before we do, I just – I mean, this is kind of crazy, but I just got an email this morning from Megan, who has a kind of a crisis. And I just wanted to ask if SpinRite would help with this. I've been trying to help her out. She apparently lost a folder within Outlook. And it had a bunch of business emails that she absolutely needed. It's not clear from her email if she lost the entire Outlook.pst. I think she did. She says, "I didn't delete anything. I didn't do anything." Is it, first of all, possible for a file just to disappear because of a hard drive error?

Steve: Well, yes, it absolutely is.

Leo: Would SpinRite recover a file if it were lost that way?

Steve: Probably. If the problem was some file system damage that was not created by a virus or a malware that is deliberately rewriting something, but if it was a sector which was no longer readable, then...

Leo: Right.

Steve: And, for example, we often see, you know, the testimonials page is full of people who were unable to get into Windows. They were like...

Leo: They couldn't even boot.

Steve: It was in a reboot loop, they couldn't boot, they couldn't do anything. Or something that they needed was broken. And so they just, you know, they ran SpinRite; and almost by, you know, almost miraculously, as they describe it, now whatever problem they had is gone because the problem was as a result of a low-level hard disk problem.

Leo: So just to make it clear, SpinRite isn't – it doesn't understand file systems. It's file system independent. It's not going to rebuild your table of contents. But if the table of

contents is damaged because of a bad sector...

Steve: Exactly.

Leo: ...or if a file is damaged because of a bad sector, it can recover that data by recovering the data, either moving it or kind of reformatting the sector underneath it to make it readable again. Is that right?

Steve: Well, that's exactly right. The way hard drives work today is different than the way they used to work. In the old days SpinRite had to understand the file system because it had to relocate the clusters of data onto a different area of the drive, and then relink the file system to refer to the relocated data. Today drives are doing sector reassignment and relocation themselves, underneath the file system.

Leo: Is that part of the ECC process, or...

Steve: Well, yeah. What actually happens is that, if you remember, when you and I demoed SpinRite on Call For Help, we could see errors occurring all the time on drives.

Leo: Hundreds of them every minute.

Steve: Yes. Yes. Drives have gotten so dense that drives are now depending upon the error correction code, ECC, just to read good sectors, not even bad sectors.

Leo: You're really at quantum level now where you're kind of guessing where the bits are and whether they're on or off...

Steve: Leo, these drives...

Leo: ...and hoping that you're right.

Steve: These drives are so dense...

Leo: It's amazing.

Steve: ...I mean, it's just insane how dense they are. So what...

Leo: I stopped buying big drives when you told me that. You know, I buy Raptors because I like the 10,000 RPMs. But I buy their 150 gigabytes. I don't buy the big drives anymore.

Steve: It's funny because my computer supplier, whenever I call him and need some drives, I

say, okay, what's the smallest...

Leo: What's the smallest.

Steve: What's the smallest drive I can still buy?

Leo: And that's because you're trying to reduce aerial density and therefore make them more reliable?

Steve: Yes. Yes.

Leo: Okay, interesting.

Steve: I mean, you know, if you can imagine that today if we needed to we could create the most reliable 100-meg drive ever seen because, you know, by comparison the 100-meg drives, the individual bits were so huge on the drive that there's no way current technology could not read them reliably.

Leo: Right, right.

Steve: But, you know, manufacturers, I mean, think about it, you know, they're motivated by profit. They want to cram as much data as they can. And so drives are reliable enough. You know, they could, you know, they could make them absolutely reliable. Or, well, by increasing the price, or by lowering the density. But they, you know, they want...

Leo: They want 500-gigabyte drives.

Steve: And we want them to be inexpensive. And so...

Leo: Cheap 500-gigabyte drives, right.

Steve: Exactly. So they make them reliable enough. So anyway, but back to ECC real quickly, what the drive does is, when it makes a correction, basically it creates an EXOR mask. And we've talked about EXORing in the past in the realm of encryption, where an EXOR mask basically flips various bits. So the error correction code, when there's a problem it creates an EXOR mask that inverts the wrong bits to make them right again. Well, the length of the mask tells the drive how bad the sector is. The longer the mask is from the first 1 bit to the last 1 bit, which is to say the first bit that needs to be fixed and the last bit that needs to be fixed, in a linear run of bits, the length of that tells the drive whether it's time to relocate this sector.

So drives can tolerate, and they do now routinely tolerate, just all the time lots of little corrections. When the correction starts getting too long, it passes a certain threshold, and the drive says, uh-oh, if I wait much longer this may become uncorrectable. And so the drive takes the sector out of service, copies the corrected data to a new sector, and internally swaps that back into place.

Since all of that happens underneath the file system, SpinRite doesn't need to worry about whether it's a Linux or a TiVo drive or a Mac or NTFS or anything. It just works with the drive at the sector level to repair things. And SpinRite's magic is it's able to recover unrecoverable data. That's always been its claim to fame. And it's able to do much more than the drive can do by itself, working with the drive. So, you know, that's a long-winded answer to your...

Leo: Well, what I'm going to do is I'm going to send Megan an email saying, you know, get SpinRite. I just wanted to verify with you my initial instinct, which was it could be, in fact, just a bad sector.

Steve: Well, and you know, for people like her, who might be using SpinRite for a specific purpose, you know, we don't have a demo because there's never been a way...

Leo: Lot of people say I want a demo of SpinRite, I know.

Steve: I know. So what we do is we have, you know, no problem with someone buying it and trying it and then writing to our sales address and saying, hey, you know, it didn't work for me, I'm going to destroy my copy, I'd like my money back.

Leo: Wow.

Steve: And so, you know, we just have no problem with that at all.

Leo: Oh, I had no idea you did that. That's really...

Steve: Oh, no questions asked, we'll just put the money back on your credit card.

Leo: Steve, you're so ethical, I can't stand it. You're going to – I don't know how you've survived in this industry for as long as you have.

Steve: Well, it's because, you know, look at that testimonials page, you know, it...

Leo: Well, you were telling me about one guy who pirated it.

Steve: Yes. In fact, we got a piece of email last week, he said...

Leo: This cracks me up.

Steve: I've got it in front of me. He said, "I must be honest and tell you I used a pirated copy of SpinRite..."

Leo: Shame.

Steve: "...a friend of mine gave me to recover..."

Leo: Because you've never copy protected it. I want to emphasize this. He's never put copy protection on SpinRite.

Steve: No. I mean, I would much rather trust our customers...

Leo: Yes, yes.

Steve: ...and not impede them in any way.

Leo: Yeah.

Steve: I mean, you know, more and more...

Leo: Good for you.

Steve: ...more and more, you know, copy protection is becoming a problem. Anyway, so he says, "...a friend of mine gave me to recover 120 gig of data on a drive which suddenly quit. As a systems admin by trade, I should know better and back up regularly. But you know how it is. Backups are always on a backburner, not something we think about until problems occur." He said, "Needless to say, SpinRite saved my bacon. I was able to recover all my data..."

Leo: Yay.

Steve: "...then move it to a secondary hard drive."

Leo: That's so great.

Steve: He says, "I was so thankful and happy, I just bought my own licensed version of SpinRite. Thank you so much for an amazing product."

Leo: I doubt there are very many people, very many companies that can say that their pirated program worked, and so people bought it. I mean, that's really nice.

Steve: Actually, Leo, and we hear this all the time.

Leo: That's really neat.

Steve: I get this kind of mail. And I mean, so I don't want to formally tell people they could, you know, loan their copies. But if anyone listening to this owns a copy of SpinRite, a friend of theirs is in trouble, I'm not going to have a problem if SpinRite could fix their system. I would ask as a courtesy that that person then considers buying SpinRite...

Leo: If it works, buy it, yeah.

Steve: ...you know, if it works for them.

Leo: Yep. You know, I wish we'd had SpinRite at the datacenter where the website is stored because you know we had a drive failure over the weekend. And it was everything I could do not to get on a plane with a copy of SpinRite and fly to Dallas. You know, I don't think the techs there were paying much attention. They didn't try to recover the drive, they just pulled it out. But there was some bad sectors, and Linux was booting read-only because of it.

Steve: Yup.

Leo: And I thought, boy, if I can only get SpinRite on there, I know I could save that drive.

Steve: That's the kind of thing it does all the time, Leo.

Leo: Fortunately we had good backups, and they gave me a free drive to replace it, so it was okay.

Steve: Cool.

Leo: Although I have to say it took them so long I'm moving to a different datacenter. And now I'm using RAID 5 with three Raptors, three 150-gigabyte Raptors.

Steve: Perfect.

Leo: And I think that will – that will protect us, yes?

Steve: Yes. In fact, that is exactly my configuration. I'm not using Raptors because I don't need that much performance. But on GRC's web servers I use RAID 5. I have three drives in a RAID 5 with a hot spare.

Leo: Right.

Steve: And so the RAID controller is able to switch over and rebuild the RAID array...

Leo: Oh, that's neat, automatically.

Steve: ...on the fly, yeah.

Leo: Oh, that's neat. Well, I think I'd still have to make a call to the network center if that happened. But at least we wouldn't lose anything. And I back up, you know, we have a network-attached storage device which I back up to and everything.

Steve: Right.

Leo: We didn't lose any data, thank goodness. Anyway, enough of – the only reason I bring this up is, if you get a chance to talk to Steve about hard drives, I mean, this is the guy who knows more about hard drives probably than any human alive, except maybe Stanley Seagate or whatever. But – or Shugart died, didn't he, Al Shugart. But...

Steve: Yup, Al.

Leo: So it's nice to get – sometimes I want to ask the – I have the privilege of asking you directly, and I do appreciate that.

You know, as long as we're talking about security, let me mention our long-time sponsor – it's kind of neat to be able to say "our long-time sponsor" – and the great guys at Astaro Corporation, makers of the Astaro Security Gateway. I've told you before that it's free for home users for noncommercial use. You could download it from Astaro.com, put it on a, you know, beige box PC, any old computer lying around, and there it is, boom, you've got a great security device. For a little bit more, I think it's something like 80 euros a year, you can add antispam, antivirus, antihacker, all of this great stuff, all open source, all high quality.

I also want to tell you, though, that the Astaro Security Gateway appliance is designed for business. If your small or medium business network needs superior protection from spam, viruses, and hackers, as well as complete VPN capabilities, intrusion protection, content filtering, and an industrial-strength firewall, all in an easy-to-use, high-performance appliance, contact Astaro, Astaro.com, or call 877-4AS-TARO to schedule your free trial of an Astaro Security Gateway appliance in your business. Astaro, proud sponsors of Security Now!. Astaro.com.

Today we're going to talk about proxies.

Steve: Yes. We talked last week about MojoPac, and I said that I wanted, you know, the plan was this week to talk about MojoPac. However, unfortunately, as I tried to communicate with them, because as I said I wanted to really understand what the technology was, I did a lot more research, and I found out that, you know, at the moment they're sort of in – they've got problems. From the looks of it, they jumped out of the gate prematurely. There was the fall DEMO was the weekend before the Podcast Expo that you and I attended...

Leo: That's right, it was in San Diego.

Steve: ...a couple weeks ago.

Leo: That's right, September 26th, I think, yeah.

Steve: Yup. And Mark Thompson was actually one of the 50 companies that was chosen to show. Well, the MojoPac company was another one, as well as a third company called U3. I mention that because U3 has something very similar. Unfortunately, the MojoPac stuff has apparently had a lot of problems. It was in beta, working in a certain way. When they took it to v1.0, they made a major change that involves basically their own copy protection scheme.

Leo: Oh, boy.

Steve: It used to run on USB 1.0 drives, which do not necessarily have to have a built-in serial number. The spec for USB 2.0 incorporates a serial number which locks their MojoPac stuff to the drive. Well, so what happened was, people who were all excited about MojoPac and using it in beta, when it went to v1.0 it broke most of what people were doing without any announcement or word. Then when people started complaining, they shut down the forums. MojoPac shut down the forums in order not to allow these complaints to be seen. Which really...

Leo: Such a case study on how to do things wrong.

Steve: It was amazing. And then they had – the beta was also locked so that you could only use it for 30 days, that is, the trial, you could only use it for 30 days or 99 boots. But it turns out the way people were using it, they would boot it many times a day, so it ended up lasting a week or two, then it would refuse to work any longer. But the e-commerce system was not online yet, so people could not buy it...

Leo: Oh, boy.

Steve: ...and could not register it. Well, a lot of that has just been fixed, just in the last couple days. I've been trying to get someone on the phone. I've sent email, I mean, even to the Whois registration email addresses, because I can't get any response from anybody. And so I just thought, okay, let's – I'm going to back off a week.

Leo: Let's not talk about this product.

Steve: We can't. We really – I think it's not quite ready for primetime.

Leo: Not ready yet.

Steve: And it does seem like there are problems with the technology because it doesn't perform a full virtualization of the system. So people have been finding that registered software

that locks itself to the computer won't – it will see a different computer when you move it to a different environment. Well, the whole idea is for it to be portable so you're able to, you know, stick it in any computer you happen to be visiting, and it brings your environment with you.

Leo: Right, right.

Steve: Well, if you are relying on the programs that you've got, and they suddenly say, oh, you're not authorized, then, you know, it defeats the purpose. So anyway, while doing all this, I found another company called U3 that has something similar and far more mature just appearing. And they also announced and released at the fall DEMO show. And then there's another company called moka5, which is an interesting, really interesting approach coming out of Stanford.

Anyway, you know, we've been talking about how virtualization technology is going to be moving forward. Well, basically these are three other virtualization solutions that I want to cover next week.

Leo: Great. Because I'm very interested, and obviously there's a lot of demand for this kind of a product. Of course, if it doesn't work, it's not worth...

Steve: It's a cool idea, the idea that you carry your environment and your applications and your data with you, you plug them into any...

Leo: Everybody wants this. This is great.

Steve: ...any Windows – well, and of course it's been enabled by the fact that we now have affordable 4-gig USB flash drives. You know, you couldn't have done it back when people were talking about 64 megs...

Leo: Right, right.

Steve: ...because contemporary applications and data won't fit there. But now that you've got multiple gig flash drives, it's like, wait a minute, this is – I mean, it inverts the whole model. So, I mean, it's very exciting. And the other thing, too, is the MojoPac guys really seem to be a little bit license heavy. For example, you can move your MojoPac between drives I think twice, and never again.

Leo: Oh, please.

Steve: And it's like – I know. I know.

Leo: Why why why why why why.

Steve: And the problem is, you know, you would like – and it's not cheap. It's $29.95 now, but it's going to be 50 bucks after it comes out of its trial phase. Well, that seems like a lot of

money for what they're doing, especially with all the limitations that they're putting on it. And it seems like a great idea. So my sense is that they are not going to be alone in this market for long, and there will likely be some free or open source solutions that are coming along. And in fact, I've already pretty much found one. So...

Leo: Oh, good. All right. Well, we'll talk about all that next week.

Steve: Yup.

Leo: This week, proxies. And you see this a lot. I see proxies all the time. You used to see it more often. You know, it's hidden away in the browser settings, and every once in a while you'd see a program that'd say, do you have a SOCKS proxy? And then people were talking about proxy servers. It doesn't seem like people talk about it much anymore. What is a proxy, and what is it used for?

Steve: Well, it's the kind of thing we want to talk about because it can be – it can enhance your privacy and anonymity. And it can also be a security vulnerability and problem.

Leo: Ah, okay.

Steve: Essentially, stepping back and looking at a definition, you know, the word "proxy" itself implies somebody acting on your behalf. You know, when you give, like, voting proxies for somebody else, you are saying, you know, you have my voting proxy. You can vote the way, you know, I instruct you to. And so they do that on your behalf. The formal definition of, for example, an Internet proxy, or a web proxy, is that your client, for example your browser, instead of connecting to the remote server, it connects to the proxy server. It makes the request of the proxy, which then turns around and, on your behalf, makes the request to the remote server. The result comes back to the proxy; it then turns around and, through the connection you have to the proxy, it returns the data to you.

This is actually being done by many ISPs. And in fact my local Cox Cable provider often proxies in a transparent fashion, that is, you as a customer of the ISP, you believe you're connecting to the remote machine. In fact, you are being transparently proxied, that is, it's not a proxy that you even see or are aware of. But you are – you're actually connecting – your connection is intercepted by the ISP. It then makes the request. The reason they do this is for caching. The ISP naturally has to pay for all the bandwidth it uses out to the Internet. But if instead it can put a large caching proxy at its boundary to the Internet, then maybe some of the things your web browser is asking for are in its cache.

So what happens is, the first person, for example, to go to Amazon.com will make a request for the Amazon page. And due to, you know, the uniqueness of the page, there are all kinds of, you know, it looks very cryptic, if you look at one of these Amazon.com pages, because they don't want their page to be given to anybody else. So what happens is you make the request to Amazon.com. The proxy intercepts it, gets the page from Amazon, and then returns it to you. As we know, we've talked about how web pages are built, that web page contains all kinds of other little pictures, you know, pictures of buttons, pictures of the user interface, all kinds of other things, which your web browser then sends out in a series of requests. Well, those are much more generic than the web page, the custom web page you've received. So it's very likely that, if anybody also behind the same proxy, in this case an ISP's whole domain, if they had recently gone to Amazon, and their browser had pulled all of those other things back through, the proxy would cache them, if it had permission to do so, and keep them locally.

Leo: That's why – that's that argument I have constantly with our advertising agency because their numbers are so low, and I think a lot of it is because what would be the number one thing you would cache? A podcast.

Steve: Yes.

Leo: If a thousand Comcast customers all ask for the same Security Now! podcast, I bet you anything it counts as one download.

Steve: Well, and it's good for two reasons. First of all...

Leo: Saves me bandwidth, that's for sure.

Steve: Well, exactly. It saves you bandwidth – okay, there's three reasons. It saves you bandwidth, it saves the ISP's bandwidth...

Leo: Right.

Steve: ...to the 'Net. Remember that they're paying for their usage...

Leo: They only get one copy.

Steve: Exactly. They get one copy...

Leo: And it's faster, too, to the local customer.

Steve: That's the third advantage, exactly, Leo, is that we're getting it from our local ISP. Even though it seems to be coming from AOL, it's actually coming from a stored copy on our ISP, so it's much faster.

Leo: Is there any way for me to tell that I'm getting a proxied copy instead of – a cached copy instead of the original copy?

Steve: There is no way from the user's standpoint in the case of a transparent proxy.

Leo: How about from the provider's standpoint?

Steve: There actually is, if they cared. This is the...

Leo: I care. I care deeply.

Steve: It's interesting. This is a problem that I had and solved with ShieldsUP because – that's how I happen to know that my local Cox provider is using a transparent proxy. Because what would happen would be, if someone at Cox, for example, went to ShieldsUP, they're using a web page that they think is coming from me, but it's actually been intercepted by the Cox transparent proxy. The IP address that I receive from – on my web server is the proxy's IP address.

Leo: Not the local user's.

Steve: Exactly. And so if I weren't, you know, a smart boy, I would be testing the security of the proxy server.

Leo: Of Cox.

Steve: Exactly, rather than the actual visitor who's coming to GRC to get their shields tested.

Leo: So you are able to get around that.

Steve: Yes. I do it two ways. First of all, SSL, as we've talked about before, is a key for avoiding proxies because, unless the proxy is explicitly configured – and we'll talk about that in a second, what it means to have an explicit proxy as opposed to a transparent proxy. But a transparent proxy cannot transparently intercept SSL because the certificate – there's no way it can give your browser a certificate that matches the site. So when you establish a ShieldsUP connection, the first thing I do is switch you into a secure SSL connection to avoid the proxy. The other thing that happens is that proxy servers will add some headers to the requests that they forward to the server, which is another – which is a nice way, Leo, that your advertising technology could know that this was actually a request coming from a proxy and not from the user, although it still wouldn't let the – it wouldn't let your server...

Leo: We wouldn't get a number, we would just...

Steve: Right.

Leo: We would know that who's requesting this is not an end user but a proxy.

Steve: Exactly.

Leo: But we would never know how many end users ended up sharing that copy.

Steve: That's true. And there's no way to know that. What happens is, in an HTTP request, there are headers at the beginning of the request that contain a bunch of different information, like, you know, what the user's client is. That's where cookies are offered along with the

request. The user's – a user agent. Well, one of the things, an optional header that could be included says – it's called – it's an X header. It says X-Forwarded-For, and then an IP address. So that's where it's saying, I am a proxy forwarding this request on behalf of this other IP. So then you'd get the IP of the proxy server and the IP of the actual user behind the request in order to see that that's what was going on. And there's another header called Via. And so that says this is coming via, you know, this server. So there are some headers that will expose that this is going on to the destination server. But over on the user side there's no way for them to know.

Leo: Well, I'm going to have to go send a note to our advertising agency and say, are you looking for these headers?

Steve: Well, now, there are some other things that are going on that could be done in order to give you reliable counts. Remember at one point, actually early in this discussion I said "if the proxy has permission to cache the content." So one way we know we can avoid caching is by using SSL. SSL used to be much more expensive. Sort of in the old days of the 'Net we would try not to use SSL because of the overhead of establishing connections. That was especially true back in the HTTP 1.0 days, where you would establish a separate connection for every single asset that you were downloading from the server. So you were just constantly establishing SSL connections that do have that crypto overhead at establishment.

Well, that's one of the reasons that the next-generation spec, the HTTP 1.1 spec, which is what virtually everyone is using today, it limits the total number of connections by default to two, so that your browser can establish two simultaneously connections to a server. But now it will stream multiple queries and responses through those two connections. So that dramatically lowers the per-asset overhead of SSL because it creates the connection once and then leaves it up until the browser finally says, okay, I'm done with this server. So it means that SSL is a far less overhead technology. And also our machines at both ends, the servers and the client users' machines are so much stronger now that SSL really doesn't represent an impediment. So that's one thing that could be done.

The other thing is that, when the server returns the request to the user's client, which may be intercepted by a transparent proxy, one of the headers in the response is an expiration header, or can be an expiration header. And there are even cache prohibition headers. It's possible for the server to say, this expires in. And it can express it as a future time or a duration. So basically, I mean, and that's really what you'd like. You'd like the server to say these GIF images that I'm serving for the page have a lifetime of a day, for example. And so what that does is, that permits not only the transparent cache, the proxy cache that may or may not be in line, but that's something that the user's own browser pays attention to. This is where you sometimes see a difference, if you refresh your page, especially at least for the IE, if you hold the Control key down when you do a refresh, sometimes you'll get a different page or a fresher page. What holding the Control key down does is it says to IE, do not use any assets in your cache. Go get fresh copies. So what...

Leo: The Shift key on Mac, for those – yeah.

Steve: Okay. So the idea is that, you know, it's a benefit for users and for the server because the server says, here's a page; and, oh, by the way, all these little, you know, page decoration things, they're good for a day. So as long as it hasn't been a day since you last asked for it, feel free to use it from your local cache, that is, your local browser cache. And what that has the effect of is it speeds up the second and third and successive pages on the same site. As long as those pages are using the same assets, your browser takes them from your own machine locally. Similarly, any transparent cache is basically obeying the same rules. It receives the assets, and it will cache them for all of the users that might be behind the proxy served by the

same ISP, as long as that thing hasn't expired.

Now, the other thing that can be done for volatile content, for example, e□commerce pages, you would hope that they're all going to be done over SSL, so they're not going to have a problem. But I've seen sites where you're only secure during the sensitive transaction involving a credit card, and then you're no longer secure. Hopefully those pages are all marked as non-cacheable. There's a header that says absolutely never cache this. The other thing that can be done – and sometimes it's done redundantly, it's what I do on my own site – is I show an expiration in 1997, which is actually a commonly used data in the past. That will also prevent anything from caching that content because it's like – it's pre-expired. It says, you know, this expired in 1997. Well, now it's 2006. And so, you know, obviously that's no good. So there are things that could be done to prevent these things from being cached.

But you're right, Leo, there's no way to know, if the content you're serving isn't pre-expired or served so that it cannot be cached, and probably is being cached, because it's such a win for ISPs.

Leo: So I have to figure – I remember that @Home kind of pioneered this. That's one of the earliest broadband ISPs. That was one of the things they were most proud of was this cached proxy servers. And I bet you that we're missing a lot of downloads from the big broadband ISPs because I'm sure Cox, Comcast, Roadrunner, they all do this.

Steve: Yes. I mean, it really is a win for them.

Leo: Yeah.

Steve: Okay. Now, what we've talked about so far has been transparent ISP caching. There's another type of proxy which is not transparent, which is something that the user, many users often in a corporate setting, they have to configure their browsers to use their corporate proxy in order to get access to the web.

Leo: That's that setting in the browser that you see.

Steve: Exactly. And here's an interesting tip. Once I went down to my office, and I think it was in the evening, because Sue was having some problems with her computer. I fired up IE and waited for a long time, I mean, more than I'm used to, for it to get going. And I thought, huh, I wonder what's going on?

Well, I poked around, and I found that she had left IE set in its default case, which is "automatically determine proxy settings." It turns out that, if you haven't disabled that, every time you launch IE you are hung for a while, while IE looks for either a DHCP or a DNS server that can point it to a file containing its proxy settings. So for anyone listening who's an IE user, if you go through the normal menu, you go to – under the Tools off the main IE menu to Internet Options, then choose the Connections tab, and then at the bottom is a LAN settings button. You should normally, unless you need to use a proxy, you should uncheck "Automatically determine the proxy settings." If you do that – I'm not sure if you have to restart IE or reboot, probably at least you'll have to restart IE – you'll be surprised how much faster Internet Explorer gets going. Because the way these proxy settings are determined automatically is that Internet Explorer sends out a request to a DHCP server.

Now, we've talked about DHCP frequently because that's – it's what's engaged when you tell

your system to obtain its IP settings. Well, it turns out that DHCP can actually provide a much wider range of information if it's asked to. So beyond just my IP address, my subnet mask, and my gateway, it's actually able to provide additional information and, in fact, in this case a URL for a server that can give the browser a file telling it how to configure itself for local proxy usage.

Leo: Oh, that's cool. I like that.

Steve: It is cool. Now, the bad news is, first of all, this was all designed by Microsoft for IE in a 1999 draft, you know, an IETF, an Internet Engineering Task Force draft. It has never been ratified. It was never adopted. It's expired. In fact, I think it expired in 1999. I think it was actually older than that. It expired; it was never renewed. And of course here we are now in 2006. So it never really took hold. The problem is, if DHCP, if no local DHCP server responds, Microsoft's IE falls back to using DNS queries. It uses your machine name, you know, within a corporate environment you might have, like, your machine name or number, then, you know, dot your division, dot your branch, dot your company.com, for example. And so what it does is, it looks for a specific DNS entry, and it goes backwards down the line, stripping off one dotted name at a time, until it gets all the way back down to your company.com, or org, or edu or whatever.

Well, so obviously that takes time, which is where all this delay is coming from as we're waiting for IE to launch. It always does that if you have it set to obtain the proxy settings automatically. The problem is, it makes a mistake if you've got more than a single name in your top level domain. You know, it works fine for, you know, GRC.com or, you know, UCI.edu, for example. It fails, though, with those URLs like co.uk. It actually thinks that co is your company name.

Leo: Oh, please.

Steve: It ends – I know. It's such a crock. It ends...

Leo: It's so easy to figure that one out.

Steve: And it ends up being a security problem because it makes a request out onto the Internet for, I think it's a file called W – oh, it's wpad.co.uk. Because it thinks you're at a company called co in the top-level domain uk.

Leo: Right.

Steve: And it turns out that that file, I mean, that machine was created, and it can give your browser bogus proxy settings, causing it to reroute all of its traffic there.

Leo: Oh, interesting.

Steve: Thank you, Microsoft.

Leo: Boy, you'd think they'd know a little bit about TLDs by now.

Steve: Yeah.

Leo: It's just not that hard.

Steve: Well, it makes a mistake. Now, I'm not sure that the latest versions haven't fixed this. Who knows what's going on. But it's definitely the case that anyone who doesn't need to use a locally configured proxy server definitely wants to go in and turn off that automatic proxy discovery technology. It's bad in the first place. And it slows down every time you use Internet Explorer and fire it up for the first time.

So what this does, essentially, is it gives your browser a URL and a port. Often it's, like, port 8080 that is used internally. And all of your browser's traffic is rerouted through your company's local proxy server. Now, that has a number of different benefits, you know, which really can be useful and enhance the experience. For example, it can be doing caching. It will often be enforcing policies. You might have your corporate IT, you know, blacklisting a whole bunch of sites that they don't want anyone going to, so it allows them to perform filtering. It also allows them to perform, you know, on-the-fly malware filtering, so that you're not going to be getting malicious scripts and code and things, by forcing the browser to go through this locally configured proxy. So it can be that this stuff is being done transparently.

But more often in a corporate setting you'll find that, you know, if you don't use their local proxy, they've got port 80 blocked at the corporate firewall, so you're not able to get out and access a server on port 80. You are forced to configure your browser to use your employer's proxy, which in the process of doing so means that all of your traffic goes through their filter for, you know, for whatever purposes. And unfortunately, it does mean that, you know, there are privacy concerns.

Leo: Right, right, right. So what should you do?

Steve: Really you have no choice at that point. You could – and we sort of talked about this – you could use, you know, like a Hamachi, or GoToMyPC...

Leo: Or VPN...

Steve: ...or a VPN tunnel, in order to get out of the corporate shield. And of course, you know, then there's a problem. If your corporation is allowing that, that's controversial because then you are explicitly avoiding whatever, you know, security policies they've put in place for the purpose of protecting themselves.

Leo: Right, right.

Steve: So that's, you know, that's of dubious value. And, you know, you could imagine people in the future are going to be, you know, blocking these sorts of things because, you know, it does open holes through the corporate security.

Leo: Right, right.

Steve: And finally, there are other proxies, I mean, just sort of in general. There are, for example, compression proxies, which cell phones and PDAs are beginning to use more and more. They will connect through a proxy. And its job is to reformat large pages to fit much smaller screen sizes. So there you're actually seeing, you know, content being redesigned on the fly in order to fulfill some specific niche need of that proxy user. And so, I mean, the whole concept of a proxy is something which is in between you and your remote connection and is probably doing some sort of filtering function. There's even a very popular local proxy that you actually run in your own machine, called Proxomitron. Proxomitron...

Leo: Oh, yeah, we've recommended that for years.

Steve: Yeah, I mean, it's a terrific...

Leo: It's the worst-looking piece of software I've ever seen, but it does get the job done.

Steve: Yes. It is – it's not – and it's extremely tech-y, and not very user friendly. But the idea is, it sets up a little local web server on your machine. You configure your web browser to use it as, you know, your own machine as the proxy server. So you make your request to Proxomitron, just like we were saying. Then it issues the request from itself in your computer. It comes back. That allows it to impose itself in between your browser and the outside world and perform whatever filtering and, you know, antispam, anti-advertising, all the different things that it's able to do on behalf of your system.

Leo: That raises an interesting question. What kind of performance hit do these proxy servers impose?

Steve: Well, now, that's a very good point, Leo. Certainly the fancier they are, the more they're doing, the more performance penalty you're going to have. And a simple caching server is very low overhead. You know, the instance of a caching server that an ISP has, as we've said, due to the fact that it will often have substantial amounts of content that are current in its cache due to somebody else having recently gone to MSNBC or MSN or Amazon or eBay or whatever, I mean, the net benefit is a huge gain because you're not loading down that remote server, the destination server, and you're getting your content back on essentially what is a LAN, you know, your own ISP's LAN is, you know...

Leo: So that's a big improvement.

Steve: ...all of the...

Leo: That speeds up everything.

Steve: Yes. And also...

Leo: But something like Proxomitron might slow it down because...

Steve: Yes. And also Proxomitron is very powerful, inasmuch as it's using regular expressions, it's using scripting, it's got typically lots of layers of things. One of the things that I've talked about and just have never gotten around to is this notion of GRC's net filter, which would be a transparent filter to do these sorts of things, but of course I'd write it in Assembly language; and, you know, I'd really...

Leo: Make it efficient.

Steve: Make it really efficient and way prettier and more user friendly. So it's on my list of things I hope to get around to one of these days.

Leo: Yeah, someday.

Steve: I know.

Leo: It's a long list, folks.

Steve: It's a growing list.

Leo: Long and growing.

Steve: And lastly, I'm going to reserve an in-depth look at one additional proxy system for its own entire episode because it's so cool. Many people are interested in it and are talking about it. And that's TOR, The Onion Router system.

Leo: Yes, we get lots of questions about TOR.

Steve: Yes. It's really cool technology. I want to give it its own show because it deserves it. And it's basically – it uses a chain of explicitly privacy-enhancing proxies in a very clever way, using essentially nested encryption that creates – the reason it's called an "onion router" is it's like layers of an onion, which are wrapped and unwrapped. And we've got all of the background from our crypto series to understand the use, its use of symmetric and asymmetric crypto keys. I mean, it's just been beautifully designed. So we're going to talk about that here before long, in a couple weeks.

Leo: Good. Well, we made this an extra-long episode because I slowed you down at the beginning. But I'm glad we talked a little bit about proxies. I know people are very interested in that, and we've seen that word around for a long time. And there are all different kinds of proxies. Some could be an advantage; some could be a security issue. And it just depends on how it's being implemented, I guess.

Steve: There actually have been cases where proxies have, by mistake, they have cached

malicious pages.

Leo: Wow.

Steve: And so after the server was taken down, the server's content had been scattered all over the 'Net in ISP transparent caching proxies which continued to infect people's systems.

Leo: That's bad. You've got to be careful what you cache.

Steve: Well, and in fact there have been search engines that have been guilty of the same thing.

Leo: Right, right.

Steve: Since search engines cache, you know, basically, you know, Google – I don't mean to single out Google because all search engines do this, or the ones that are doing caching do – they will keep copies of pages. They don't know that it's a malicious page. But they'll keep it in their local cache. And if you ask, they'll, you know, serve it to you.

Leo: Right. It's there.

Steve: Yup.

Leo: Yup. Although I would imagine, certainly in this day and age, that Google would spend some effort trying to sanitize that stuff.

Steve: Well, of course, as we know, Google is now explicitly using a list of known malicious pages and warning people, if they click on a link – I actually got that warning a couple days ago when I was poking around.

Leo: Oh, really.

Steve: Yeah. It'll say, this is a potentially malicious page. Are you sure you want to view it?

Leo: Oh, that's really interesting. I've never seen that.

Steve: It's really a cool service.

Leo: Yeah, that's a good idea. All right, Steverino. I appreciate your time, as always. We remind people that SpinRite is the ultimate disk recovery and maintenance utility and available from Steve's site at GRC.com. We also thank the folks at Dell for providing the

support for this podcast. They've been a great sponsor for the last couple of months. And every week we add new picks to the Leo's Picks Page at TWiT.tv/dell. If you're getting ready for Vista, you're ready to buy a new computer, you'd just like to see what's out there in terms of laptops and desktops, we've got laptops and desktops and other things all on the page. And we'd appreciate it if, even if you don't buy those items, if you click through there, we get credit for anything else you do buy at TWiT.tv/dell. Thanks to Dell for their support.

Steve Gibson, I think we've wrapped up another thrilling, gripping episode of Security Now!. We're going to call it Security and Hard Drive Maintenance Now!. A new name for the show. I apologize for bringing that stuff up. I just, you know, when I can have you to ask questions, it really is a valuable opportunity for me. And I know people are interested in these subjects, so...

Steve: Yeah, I've actually had people want to spend more time talking about SpinRite, which I – I mean, I want to keep it in people's minds because...

Leo: Even just how drives work would be of interest. It's not security, but you know so much. I'm sure in order to write SpinRite you had to really learn the details of drive operation right down to the nitty-gritty.

Steve: It's my life.

Leo: Is it, you know, it sounds like it's changed a little bit since when you first started writing SpinRite.

Steve: Oh, yeah. I mean, remember that, you know, the original catalyst for SpinRite was interleave optimization.

Leo: Oh, yes, I remember that.

Steve: Yes.

Leo: You'd set the interleave size when you first formatted a drive. Oh, those dark days are gone, thank goodness.

Steve: The original concept was to do a nondestructive, low-level reformat.

Leo: Right, I remember that.

Steve: No one had ever heard of it. And in fact, when Peter Norton invited me up to Santa Monica because he wanted to buy SpinRite from me, he said, you know, Steve – and of course he was famous for the Norton utilities.

Leo: Right.

Steve: And their own reviews, their surveys of what their customers of the Norton utilities wanted, they said they wanted SpinRite. And so he had me up for lunch to, you know, what's it going to cost to buy SpinRite from you? And of course not selling it to Peter was the best decision I ever made in my life. I mean, literally.

Leo: Is that true? Why is that?

Steve: Oh, because, you know, that was in the early days of SpinRite. And I've been, you know, basically thriving from SpinRite, and SpinRite's been thriving from me, you know, for all these years. If I had sold it to them...

Leo: That'd have been it, one time, be out, and that's it.

Steve: ...I wouldn't have it, but then neither would they. Actually, they ended up cloning it. And I'm being polite because I know this is not Peter's fault. But we found our code in their version which they called Calibrate.

Leo: Uh-oh. Uh-oh.

Steve: After I declined to sell it...

Leo: I remember Calibrate, yeah.

Steve: ...they came out with Calibrate. But it ended up, we ended up winning from that, too, because since they hadn't really written it, they couldn't support it. And so their customers would call, and their own support people ended up saying, well, why don't you just get a copy of SpinRite?

Leo: That's what we stole it from anyway.

Steve: So we ended up getting referrals from them.

Leo: Oh, that's hysterical.

Steve: But Peter said – he said, "You know, Steve, when I heard that you were doing a nondestructive, low-level reformat," he said, "I just shook my head and kept waiting to see a mushroom cloud coming up from, you know, south of me in Southern California." He says, you know, "How could you possibly do that safely?" And I said, "Well, that's, you know, that's my magic, so."

Leo: And of course the mushroom cloud eventually came from the other side of the valley. But that's another story entirely.

Steve: Exactly.

Leo: For another day. Steve Gibson, a pleasure, as always. GRC.com for show notes, the 16KB versions, and of course Elaine's great transcripts. GRC.com. And we will see you next Thursday, and maybe we'll talk about MojoPac and...