The Strange Tale of the DENIAL OF SERVICE Attacks Against GRC.COM

by Steve Gibson, Gibson Research Corporation

Page last modified: Jun 02, 2001 at 22:48

Nothing more than the whim of a 13-year old hacker is required to knock any user, site, or server right off the Internet.

I believe you will be as fascinated and concerned as I am by the findings of my post-attack forensic analysis, and the results of my subsequent infiltration into the networks and technologies being used by some of the Internet's most active hackers.

An Open Letter to the Internet's Hackers

What Happened?

On the evening of May 4th, 2001, GRC.COM suddenly dropped off the Internet:

GRC.COM Bandwidth Utilization

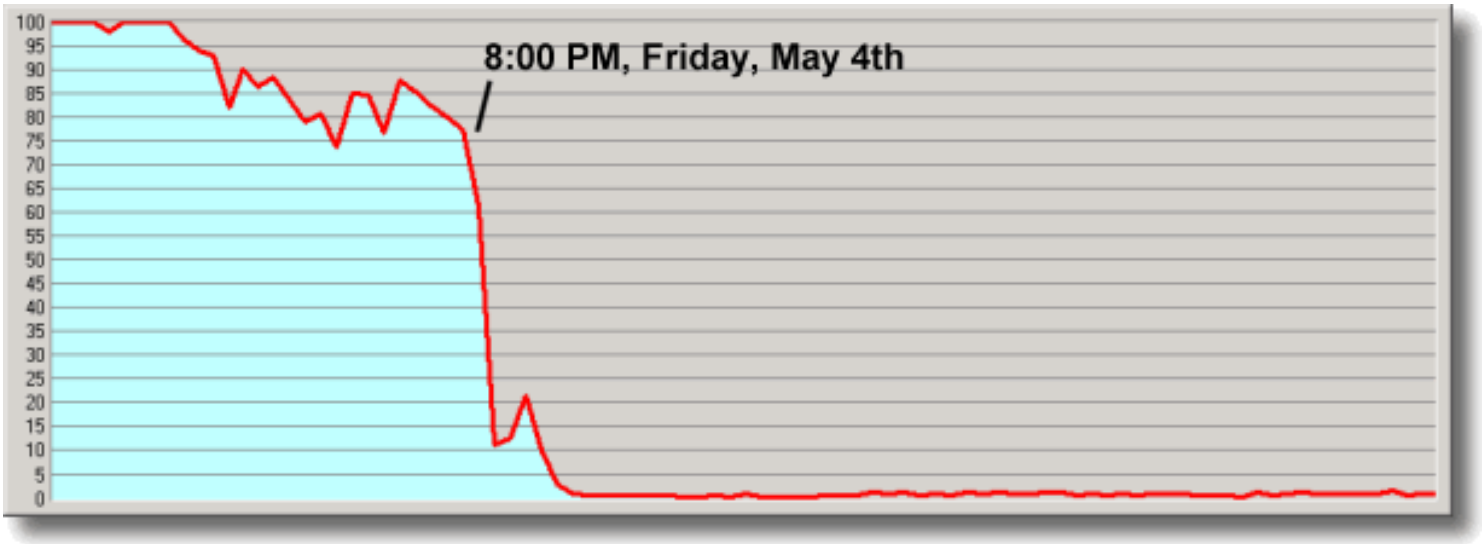

Stemming the Flood with Our ISP

Within a minute of the start of the first attack it was clear that we were experiencing a "packet flooding" attack of some sort. A quick query of our Cisco router showed that both of our two T1 trunk interfaces to the Internet were receiving some sort of traffic at their maximum 1.54 megabit rate, while our outbound traffic had fallen to nearly zero, presumably because valid inbound traffic was no longer able to reach our server. We found ourselves in the situation that coined the term: **Our site's users were being denied our services.**

I had two priorities: I wanted to learn about the attack and I wanted to get us back online.

I immediately reconfigured our network to capture the packet traffic in real time and began logging the attack. Dipping a thimble into the flood, I analyzed a tiny sample and saw that huge UDP packets — **aimed at the bogus port "666" of grc.com** — had been fragmented during their travel across the Internet, resulting in a blizzard of millions of 1500-byte IP packets. Mixed into this appeared to be ICMP debris from large-packet ping commands.

We were drowning in a flood of malicious traffic and valid traffic was unable to compete with the torrent.

At our end of our T1 trunks, our local router and firewall had no trouble analyzing and discarding the nonsense, so none of our machines were adversely affected. But it was clear that this attack was not attempting to upset our machines, **it was a simple brute force flood**, intended to consume all of the bandwidth of our connection to the Internet **. . .** and at that it was succeeding all too well.

Gibson Research Corporation is connected to the Internet by a pair of T1 trunks. They provide a total of 3.08 megabits of bandwidth in each direction (1.54 megabits each), which is ample for our daily needs.

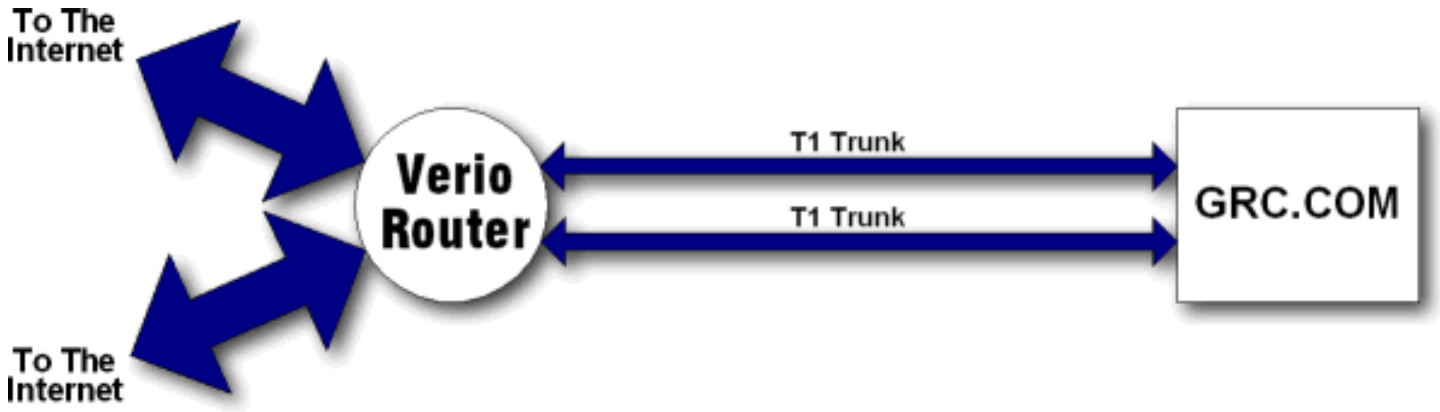

As you can see from the schematic diagram above, the Verio (our ISP) router that supplies our T1 trunks enjoys two massive 100 megabit connections to the Internet. But from there all of the traffic bound for us must be funnelled through our two T1 trunks. Therefore, in order for the congestion of our T1's to be relieved of the malicious traffic, the "bad packets" had to be filtered **before they left Verio's router**. In this way, the packet flood would be stopped at a high-bandwidth point — upstream of the T1 choke point — thus allowing the "good packets" to slip past the bad packets and cross our T1's in peace.

It took us a while . . .

Due to a comedy of changed telephone area-codes, mistyped eMail addresses, slept-through pager beeps, the attack's Friday night timing, and Verio continually insisting that I be processed through the regular channels of "the system" (which I kept explaining did not seem to be working), **seventeen hours passed before I was able to get a competent Verio engineer on the other end of the phone.** But once I had a competent engineer on the phone, and armed with my analysis of the attack pattern **. . .**

In two minutes we applied "brute force" filters to the Verio router, shutting down all UDP and ICMP traffic ... and GRC.COM instantly popped back onto the Internet.

That part was pretty cool. We were still very much under attack, but **because the attack was prone to filtering** (thank goodness) we were able to have Verio's router "weed out" the bad packets and return us to almost normal operation.

An Attack "Prone to Filtering" ?

Yes. Fortunately — as we'll see below — the attacking machines were all security-compromised Windows-based PC's. In a fluke of laziness (or good judgement?) that has saved the Internet from untold levels of disaster, Microsoft's engineers never fully implemented the complete "Unix Sockets" specification in any of the previous version of Windows. (Windows 2000 has it.) As a consequence, Windows machines (compared to Unix machines) are blessedly limited in their ability to generate deliberately invalid Internet packets.

It is impossible for an application running under any version of Windows 3.x/95/98/ME or NT to "spoof" its source IP or generate malicious TCP packets such as SYN or ACK floods.

As a result, Internet security experts know that non-spoofing Internet attacks are almost certainly being generated by Windows-based PC's. Forging the IP address of an attacking machine (spoofing) is such a trivial thing to do under any of the various UNIX-like operating systems, and it is so effective in hiding the attacking machines, that no hacker would pass up the opportunity if it were available.

It is incredibly fortuitous for the Internet that the massive population of Windows-based machines has never enjoyed this complete "Unix Sockets" support which is so prone to abuse. But the very bad news is **. . .**

This has horribly changed for the worse

with the release of Windows 2000 and the pending release of Windows XP.

For no good reason whatsoever, Microsoft has equipped Windows 2000 and XP with the ability FOR ANY APPLICATION to generate incredibly malicious Internet traffic, including spoofed source IP's and SYN-flooding full scale Denial of Service (DoS) attacks! (See my WinXP & DoS Page.)

While I was conducting research into the hacker world following these DoS attacks, I encountered evidence — in attack-tool source code — that malicious hackers are already fully aware of the massive malicious power of the new versions of Windows and are waiting impatiently for the "home version" of Windows XP to arrive in the homes of millions of less clueful end users.

When those insecure and maliciously potent Windows XP machines are mated to high-bandwidth Internet connections, we are going to experience an escalation of Internet terrorism the likes of which has never been seen before.

If I fail in my mission to convince Microsoft to remove this from Windows XP, the historical problems with Internet attacks promise to pale in comparison to what will begin happening as Windows XP is deployed next year.

Thanks to the fact that the fleet of attacking machines were Windows PC's, they were unable to send TCP SYN packets to our port 80 (which would have crippled us completely), and were only able to flood us with UDP and ICMP packets (which we could temporarily ignore).

Working with our ISP we were able to filter our receipt of the malicious packets before they were able to reach our T1 trunks. We were able to continue offering our TCP-based services (Web/FTP/News) even while under continuing attack.

The Attack Profile

We know **what** the malicious packets were, and we will soon see (below) exactly how they were generated. But we haven't yet seen where they all came from. During the seventeen hours of the first attack (we were subsequently subjected to several more attacks) we captured 16.1 gigabytes of packet log data. After selecting UDP packets aimed at port 666 **. . .**

I determined that we had been attacked by **474** Windows PC's.

This was a classic "Distributed" Denial of Service (DDoS) attack generated by the coordinated efforts of many hundreds of individual PC's.

- Where do these machines reside?
- Who owns them?
- Who are their ISP's?
- What sort of users are running Windows PC's infested with potent Internet attack Zombies?

A determination of the network domains hosting the attacking machines revealed the following, hardly surprising, cast of Internet end user service providers:

104	home.com	5	inreach.net	3	voyager.net
51	rr.com	5	telus.net	3	lvcm.com
20	aol.com	5	gtei.net	3	co.uk
20	mediaone.net	4	tpo.fi	2	cdsnet.net
17	uu.net	4	rcn.com	2	enter.net
14	btinternet.com	4	isoc.net	2	cgocable.net
14	shawcable.net	4	uswest.net	2	knology.net
14	optonline.net	3	dialsprint.net	2	com.au
14	ne.jp	3	a2000.nl	2	fuse.net
9	chello.nl	3	grics.net	2	lrun.com
9	ntl.com	3	linkline.com	2	dialin.net
8	videotron.ca	3	eticomm.net	2	bellsouth.net
7	ad.jp	3	prestige.net	2	psnw.com
7	psi.net	3	warwick.net	2	pacificnet.net
6	uk.com	3	supernet.com	2	tds.net

Domains hosting two or more security-compromised, attack Zombie, Windows PC's

(The balance of the 474 Windows PC's not represented above were scattered across a multitude of domains, one machine apiece.)

It probably comes as no surprise that the top two U.S. residential cable-modem Internet service providers — **@Home** and **Road Runner** — provide Internet connectivity to the host machines most often sought by malicious hackers for the installation of bandwidth flooding Zombie attack Trojans.

While I was monitoring several online hacker hangouts (with the aid of custom spy-bots I created for the purpose — more on that below), I often overheard hackers referring to various lists of "cable Bots" and saying things like "Heh, but how many of his Bots are cable?"

It is clear that the "cable Bot" — a remote control Zombie program installed on a high bandwidth, usually on, Windows machine — has become a highly sought-after resource among malicious "Zombie/Bot running" Internet hackers.

Five Additional Attacks

After the first May 4th attack, we were left alone for eight days. As far as I knew then, the attack had been a one time event like so many of the denial of service attacks you hear about. But then, shortly before midnight Saturday May 12th, we were hit with a repeat of the previous attack lasting eight hours into May 13th.

Here is a summary of the attacks and our actions:

- **May 4th — First attack**, 17 hours before we were able to filter it at our ISP. The attack continued unabated behind the ISP's filters.
- **May 13th — Second attack**, identical to the first and using the same attacking Windows machines. Our communication with Verio was still problematical (at best), so we were knocked off the Internet for eight hours until Verio re-established our previous filters. Unfortunately, I was beginning to learn that Verio's advertised and promised 24/7 service was really only 8/5.
- **May 14th — Attacks 3a and 3b**. Since Verio was still blocking malicious traffic to the grc.com server, this third attack was retargeted at the IP of our firewall. This was one machine closer to the Internet and still on our side of our T1 trunks. So the malicious traffic was again crossing our T1's and we were promptly knocked off the Internet.

  By this time I had the home phone numbers for key members of Verio's engineering staff (<grin>) so we quickly changed the router's filter to block the firewall's IP address and grc.com popped back onto the Internet.

  However, the hackers were apparently watching our success at blocking their attacks. So two hours later the attack resumed, this time aimed at one of the two T1 interfaces of our Cisco router. This again knocked us off the Internet because malicious traffic was again crossing our bandwidth-limited T1's to our local router. Back on the phone to Verio, we decided to completely shut down that T1. GRC.COM limped back onto the Internet running on a single T1.
- **May 15th — Attack 4**. This one started earlier in the evening (at 5:00 PM PST) and knocked us off the Internet for six and a half hours (until 11:30 PM PST) before fading of its own accord. Due to bugs in Cisco's v12.0 IOS routing software (which we fully characterized and worked around several days later) we were unable to comprehensively filter this attack. Rather than engaging in another night of cat & mouse "guess the IP" as we had the night before, I decided to remain off the Internet, collect attack logging data, and take the opportunity to defragment our server's hard drives while weathering the storm.

As you might imagine, these attacks had gone from initially being interesting — from an Internet Security standpoint — to becoming a significant distraction and annoyance. So, on Wednesday May 16th, I got my local (really good) Verio engineer on the phone and we went to work methodically designing and testing a comprehensive set of Cisco router filters so that Verio's router could protect the entire grc.com domain including its T1 interfaces.

By this time, I had assembled an exact profile of the malicious traffic being generated during these attacks. Thanks to the fact that they were sourced from security-compromised Windows machines, they were subject to filtering. During this process we discovered, confirmed, and worked around the several significant packet filtering bugs in Cisco's v12.0 IOS which had caused our previous headaches.

Finally, ducking down behind these new filters, we held our breath and headed into another evening of malicious Internet attacks **. . .**

- **May 16th — Attack 5**. Being as prepared as we now were paid off completely. The grc.com domain never felt the massive attack being waged on the other side of Verio's router. The next day I asked my local Verio engineer how many "UDP/666" malicious packets had been stopped by our filters:

  12,248,097

  The attacking Windows machines generate maximum-size 64k byte UDP packets, but only the first 1500 byte "fragment" of each packet carries the packet's port "666" destination. Therefore, for every identified "666" packet blocked, approximately 43 additional maximum-size "packet fragments" were also blocked. We therefore estimate that our filters running in Verio's router blocked at least **538,916,268** malicious packets that night.

- **May 17/18/19/20th — Attack 6**. The exact dates and times are unknown because we were completely shielded by the configuration of Verio's router. But when we checked the router's "UDP/666" hit counter on the morning of Monday May 21st we found that the blocked "666" count had jumped from its previous value of 12,248,097 to a whopping 54,528,114 packets, leading us to conclude that the filters had weathered, by this time, at least :

  2,399,237,016 malicious packets.

  — nearly 2.4 BILLION malicious packets.

If the attacking machines had been running Windows 2000 or the home-targeted version of Windows XP, as they certainly will be next year, we would have been utterly defenseless and simply forced off the Internet.

This is what anyone on the Internet can soon expect.

(See my WinXP & DoS Page for more information about this.)

"Wicked" Speaks

On May 15th, after weathering our fourth DDoS attack, the following newsgroup posting, claiming responsibility and credit for the multiple attacks against GRC.COM, appeared on our news server:

```
hi, its me, wicked, im the one nailing the server with
udp and icmp packets, nice sisco router, btw im 13, its
a new addition, nothin tracert cant handle, and ur on a
t3.....so up ur connection foo, we will just keep comin
at you, u cant stop us "script kiddies" because we are
better than you, plain and simple.
```

Welcome to the brave new world of the 13 year-old Internet terrorist.

I immediately and politely replied to "Wicked" via the newsgroup. I invited him to write to me via eMail through any anonymous channel of his choosing. I wanted to understand what had precipitated these multiple attacks, and I wanted to see what I could do to get them to stop.

"Wicked" soon wrote to me through an eMail account at YAHOO.COM:

```
yo, u might not thing of this as anyomous, but its
not real info, its a stolen earthlink, so its good,
now, to speak of the implemented attacks, yeah its
me, and the reason me and my 2 other contributers
do this is because in a previous post you call us
"script kiddies", atleast so i was told, so, i
teamed up with them and i knock the hell out of
your cicso router, and....im building up more bots,
no, not sub seven lame ass script trojans, i made
my own, and it seems quite effective does it not?
seems to me that ur backbone has trouble handling
the crap sent at it, go ahead and drop icmp pings,
u still need to say "NO" to them so it still takes
bandwith, thats where tracert comes in, to find the
t3 box ur on, nice, i see u stop it as-of today, :)
good for you, now ill find ways around it and we
can keep playing these games, i find it very fun,
shout out to hellfirez and drgreen, and yeah the
hellfirez from subseven, hes a friend and he isnt
a script kiddie u stupid fuck...now, if u wish to
talk to me in person, hows irc??? talk to WkD, the
nick wicked was taken, good luck :)
```

My reply to this note from "Wicked" carefully explained that he really had me all wrong. I pointed to my "Acknowledgement of Debt to the World's Hackers" at the bottom of my "NanoProbe" page. I explained that while I did feel there was a distinction between an elite hacker and a script kiddie, I was someone who always took pains to be respectful of others' egos (when possible), and that I was unlikely — unless provoked — to casually refer to anyone using a derogatory term. I told him that while I was aware of a dispute that had erupted several weeks before in one of our newsgroups, and reportedly involved his friends "HeLLfiReZ" and "DrGreen", **I** had neither read nor participated in any of that conflict.

"Wicked" replied that perhaps he had "misjudged me", since he was only going by word of mouth. He volunteered to speak with his friends and call off the attacks. He promised that there would be no further attacks from then on **. . .** after which he attacked us on the evening of May 16th, saying afterwards:

```
is there another way i can reach you that is
secure, (i just ddosed you, i aint stupid, im
betting first chance ud tracert me and call fbi)
you seem like an interesting person to talk to
```

"i just ddosed you ..." Indeed.

Fortunately, that was the first night of our new and (so far) impregnable router filters, so we felt nothing across our T1's while Verio's router counted and discarded nearly five hundred and thirty-nine million (538,916,268) malicious bandwidth-consuming attack packets.

From my dialog with "Wicked", I saw that these repeated attacks were "fun" for him. He was like a child pulling the legs off a spider to see what it would do, watching it flail and attempt to get away from its tormentor. And, as we have seen, he experiences absolutely no remorse and has no regard for any damage being done as a consequence. He believes that he can not and will not be caught. Hiding behind the anonymity created by the Internet's trusting technology, he exhibits no social conscience.

I hope it is becoming clear to everyone reading this, that we can not have a stable Internet economy while 13 year-old children are free to deny arbitrary Internet services with impunity.

I wanted these attacks to stop, but I was certainly in no position to make any "parental" demands of "Wicked". While we were essentially functional, hiding behind our router filters, we could not remain behind them forever. We were unable to send and receive "ping", "trace route" and UDP fragments — all crucial requirements for full Internet function. In the long term this would pose serious problems for the delivery of GRC's Internet security testing services.

I had to find a better solution.

Earthlink Turns a Blind Eye.

I wasn't yet sure what I wanted to do about this 13 year-old problem, but I felt that I ought to broaden my options. "Wicked's" several postings to our newsgroups, and his eMail to me through his yahoo.com account, were all originated from the same small IP address range corresponding to the small ISP Genuity, BBN Planet, in Kenosha, Wisconsin — an Earthlink reseller. This correlated perfectly with "Wicked's" claim to be using a "stolen earthlink".

Since I had collected the exact dates, times, and IPs for each of "Wicked's" several dial-up connections, I felt that perhaps Earthlink could preserve the access and phone records in case the FBI might need them later. If his home phone number could be determined, we could identify him. I knew that Earthlink would never reveal such information to me, but I just wanted them to preserve the evidence against the possibility of future need.

Two months before this, Earthlink's privacy officer, Les Seagraves, and I met and formed a good relationship during our quest to understand the peculiar Earthlink Browser Tag. Unfortunately, Les' voicemail explained that he would be out of town through the end of the month. So I got the name of Earthlink's director of corporate communications, Kurt Rahn, from a well-placed press contact of mine. Kurt was prompt with eMail, and he made lots of motivated-sounding noises, but nothing more ever happened. After waiting hopefully for several days, I finally spoke to Kurt on the phone and allowed myself to sound a bit perturbed. It had no discernible effect. His many promises to have Earthlink's security people get in touch with me never resulted in a single contact from anyone.

Hackers take note: Earthlink appears to be a safe haven for your operations. From everything I have seen, Earthlink couldn't care less WHAT you do, so long as someone pays the bill.

Further note: The day after I wrote this, Les Seagraves returned from his trip and immediately returned my original voicemail message. Les was sincerely apologetic and wonderful when I explained the situation. So I feel self-conscious over being as harsh about Earthlink's response as I have been here. But what I wrote is exactly what happened, and I don't know how else we will ever get ISP's to spend some money, and get involved in security issues, unless we begin holding them accountable for their inaction.

The lights may be on, but nobody's @Home

I have recorded the IPs and account numbers of more than 100 @home subscribers who have security-compromised Windows machines currently running active Trojan attack Zombies. As we will see below, each of those machines also receives a complimentary copy of the latest version (v2.21) of the incredibly invasive Sub7Server Trojan. This grants the hacker who is controlling the Zombie — the "Zombie-master" — absolute control over his victims' machines. Among the many invasions the Sub7Server Trojan enables is monitoring every keystroke for the

purpose of capturing online passwords, credit card numbers, eBanking passwords and you name it.

Now, you might think that this would be significant to @home's chief of security, Todd Welch, but it isn't. I tried to talk to him on the phone, leaving a detailed voicemail describing the situation, but I was shuffled off into the system and asked to eMail the IP's to "abuse@home.com". Refusing to have the machine IP's disappear and never to know what, if anything, had been done, I called back the next day and got Todd on the phone. I have no idea why, but he didn't sound at all happy to be talking with me. It was as if he wished this problem would just go away — or that at least, I would.

I explained that many of the compromised and Zombie-infected @home machines were showing a machine name of ***.sfba.home.com**, which I presumed, and he reluctantly confirmed, stood for "San Francisco Bay Area". Since @home is in Redwood City on the Bay Area Peninsula, I thought that perhaps I could fly up to their offices, then he and I could make a few house calls on some Bay Area Zombie-infected @home subscribers.

I was itching to get my hands on one of those nasty nightmares that had been plaguing us for the previous two weeks so that I could take it apart and figure out what made it tick.

I told Todd that after I had dissected a Zombie, I might be able to come up with a way for @home to scan their network to find all of them. It turns out that **I have found a way**, but again, Todd and @home couldn't be bothered. He declined all cooperation of any sort, curtly adding that they work with the FBI, and no one else. As we will see next, this is a policy in dire need of change. Nice as it sounds on the surface, the realities of Federal government involvement mean that most of the time Todd and @home **. . .** do nothing.

Okay, time to call the Feds . . .

My Conversations with the FBI

In stark contrast to my frustrating conversations with commercial Internet service providers, **I had two terrific dialogs with the FBI's top cybercrime agents.** One of the agents, based in Hayward, California, oversees much of the Western region and is a frequent visitor to grc.com. (With tongue in cheek, he proudly informed me that his "Shields were UP!") The other agent, based right here in my home city, is one of those working on the case of the Russian hacker extortion ring, for which I had recently created the PatchWork utility. He said that he carries a copy of PatchWork around with him to quickly demonstrate how vulnerable Windows servers are due to missing security patches. All of this made our introductions much simpler and smoother.

Both FBI guys said similar things:

- They explained that until $5,000 of damage had been done, no crime had even been committed. That's the law. And due to the peculiar nature of GRC.COM's business model (such as it is :), these attacks were stirring up interest in my forthcoming research and it wasn't even clear that we were going to be economically damaged in any way.
- Secondly, they said that even if we did manage to meet the $5,000 minimum required for "Wicked's" activities to qualify as criminal, their staffs were overloaded and swamped with cases involving companies that had lost huge sums of money to Internet crime. Furthermore, since the cost of an FBI prosecution was in the neighborhood of $200,000, they needed to prioritize their cases based upon prosecuting criminals who were responsible for causing large dollar losses. "Wicked's" attacks, no matter how annoying, failed to qualify.
- And finally, they said that since "Wicked" was only 13 years old, nothing much would happen to him, even if the preponderance of evidence demonstrated that he was behind these attacks. They said that a couple of agents might go out to his home and have a talk with his parents, but in this country his youth was an impenetrable shield. This, of course, further discouraged the costs which would be incurred through any investigation.

Contrary to what you might presume, I did not regard any of this as particularly bad news. I felt that I should do what I could do in the legal arena, because I should. But I really didn't have any desire to be responsible for putting a 13 year-old behind bars. I have since told "Wicked" that if he doesn't wise up, in five years his "youthful offender shield" is going to dissolve and he could find himself in some serious trouble. He says that he was already in trouble with the FBI when he was eight — for hacking government servers. His computer was taken away until he was ten, then he was carefully monitored for another year until he was eleven. But now he's right back at it. <sigh>

The local FBI agent did ask for the IP's and domain names of the 474 machines that participated in the May 4th attack. I forwarded the information to him immediately, so now they have that stuff.

But I was still no closer to having any REAL answers . . .

Getting My Hands on a Zombie

An angry phone call from the head of Information Services at Texas A&M University gave me an idea. In some initial media coverage after our first attack, I had mentioned that a single machine at Texas A&M University (tamu.edu) had been responsible for generating about twenty times more malicious traffic than any other. The I.S. guy was pissed off that I hadn't contacted them immediately, and I suddenly felt like a total fool: I could have contacted him and we could have probably tracked down the infected machine. I immediately volunteered to fly to Texas and begged him not to destroy any evidence until I got there. But it turned out that the machine was within the student residential network (resnet.tamu.edu) — the student dormitories — and that since the May 4th attack, the University had closed down for late spring break. Drat!

But that gave me an idea: Although the big ISP's are apparently so big that they no longer need to care about their customers, the smaller users — like TAMU — might be much more receptive. So I sent out a mass of eMail to every smaller ISP and domain

administrator of the infected attacking machines. I explained the situation and asking for their help.

Someone discovered the Zombie and sent it to us!

I will never know whom to thank because they dropped the Zombie into our anonymous web-based Spyware drop-box. But it was all I needed to learn how these Zombie's operate and then infiltrate the Zombie High-Command **. . .**

The Anatomy of a Windows Attack Zombie

The Zombie program I received was named "**rundIl.exe**". (Note the capital "I" in the filename.) This struck me as significant since "rundll.exe" is a frequently used and often seen component of Windows systems. Changing the first lower-case 'l' (el) to an upper-case 'I' completely hides the difference under Windows 9x systems because the font used by the Windows registry renders those two characters as a featureless vertical bar. Anyone inspecting the Windows registry for suspicious files will see: "rundll.exe" and miss the fact that it's actually "`rundIl.exe`". Clever.

My inspection of the 15,904 byte Zombie program quickly revealed it to be an IRC (Internet Relay Chat) client. So I decided to sacrifice a PC to the Zombie by deliberately infecting it while keeping it under observation with a packet sniffer running on an adjacent machine. I freshly reformatted a laptop, installed a completely clean copy of Microsoft Windows, named the machine "Sitting Duck" **. . .** and turned it on.

The Zombie immediately connected with a remote, pre-programmed, IRC chat server. It then joined a secret and password key-protected channel on that server . . . and waited for instructions.

It didn't have long to wait.

I watched in fascination as many other Zombies — hundreds of others — arrived and departed the secret "Zombie meeting grounds" of the IRC server.

Somewhere, Windows users were innocently turning on their PC's. Lacking any effective personal firewall security (we will see later that BlackICE Defender provides no protection), the Zombies running secretly and silently inside those machines were connecting to this IRC server. They maintained persistent connections for the duration of that PC's access to the Internet. The Zombie and its master don't care whether the machine is cable-connected, DSL, or dial-up — though higher-speed connections are always preferred, as are machines that tend to be "on" most of the time. After all, you just never know when you're going to need to go attack someone.

While I was watching this sad drama, **suddenly and with no warning everything went crazy:** The packet sniffer's packet display became a blur as its scrollbar "thumb" rapidly shrunk to its minimum size. Thousands of packets were being logged per

second! Since I was nervous during this first incursion into hacker territory, my first thought was that I had somehow already been discovered, and my little "Sitting Duck" laptop was under attack.

But the cable-modem I was using to guarantee my anonymity revealed the truth: The RECEIVE light was dark, but the **TRANSMIT light was ON SOLID!**

I immediately shut down the Zombie-infected PC and scrolled the packet log back before the beginning of the attack. I found the command that the Zombie running in my laptop had received just before all hell broke loose **. . .**

My laptop had participated in an all-out Denial of Service attack against a machine in Finland!

Yikes! This was unacceptable. I wanted to keep active Zombies running here so that I could study their behavior, but I could not have them participating in Internet Denial of Service attacks. So I hacked the Zombie to kill its ability to send damaging packets.

From that point on I ran only "Attack-Neutered Mutant Zombies"

Later that night I received another surprise . . .

The Arrival of the Sub7Server Trojan

Not content to simply have a fleet of willing attack Zombies, "Zombie-masters" routinely download the latest version of the rapidly evolving Sub7Server Trojan into all of their "client" machines. Since Zombies are continually joining and leaving the central IRC server, Sub7 downloads are generally performed on the pool of currently connected machines several times per day. (After all, you certainly wouldn't want an old version of the Sub7 Trojan running in any of your Zombied machines.)

The Sub7Server Trojan is massively invasive. It has been designed to give its master virtually complete control over the compromised PC. This includes complete file system inventorying and file access, and real-time keyboard keystroke logging. Any user with Sub7 in their machine might as well have the hacker standing right next to them watching every move they make while using the computer.

Although it was not the focus of my research, examining the latest version of Sub7 was an education in itself. Sub7 is downloaded by the Zombie from a "free web pages" web server as a single file named "HardcoreS7.exe". When this executable file is run it breaks apart into two randomly-named files so that it is not possible to search by filename. The original "HardcoreS7.exe" file is then deleted. Sub7 insinuates itself into Windows in a few clever ways. It installs in the seldom used "run=" line of the deprecated WIN.INI file. It also installs under the "Run" key of the registry, and it inserts a much smaller 10k "runner" into the Windows Shell "command/open" key. All of this pretty much guarantees that Sub7 will keep running inside the system. It's difficult to shake it loose.

When Sub7 awakens inside a machine it begins listening on the random port it chose

during its installation. This means that security scanners can no longer scan for "the Sub7 port" since there isn't one. After setting itself up for business, Sub7 sends out notifications to hackers of its availability through two different channels: It joins a special Sub7 IRC chat server where it posts a notice listing all the details required for its connection and use:

Sub7Server v.BoNuS 2.1
installed on port: 27374,
ip: 293.492.59.748,
victim: Frog's Legs,
password: ribbit

At the same time, identical information is posted to a newsgroup server through a web server CGI script. Hackers no longer need to "scan" the Internet for vulnerable machines running Sub7 **. . .** each copy of Sub7 phones home providing complete connection details.

I Needed My Own Stealth Spy-Bots

Untangling and extracting the meaning from the packet capture dialog of the average attack-neutered mutant Zombie is never fun. Moreover, separating the good stuff from the noise just adds to the burden. So I soon realized that I needed to create my own "Zombie simulators" that would logon to IRC chat channels, hob-nob with the Zombie locals and their masters, while logging everything that transpired and automatically alerting me of anything important.

So I downloaded a copy of the Internet RFC 1459 for Internet Relay Chat (IRC) Protocol and figured out how IRC works.

By the time I was done, I had written a handful of application-specific tools for infiltrating and spying on hacker and Sub7 Trojan IRC channels. The tools chronicled dialogs, and captured references to other server, hacker and Zombie channels. They automatically spawned new instances to begin monitoring newly discovered servers. They snagged passing URLs and quickly downloaded anything that was referenced. I even got quite fancy and built a Markov-chain finite-state statistical dialog modeller. It monitored the flow of IRC channel nicknames and automated the process of determining who was talking to whom, and who were the "bosses" who commanded the most power and respect.

I learned an amazing amount about this bizarre world of zombie-running hackers. During the process I witnessed a truly disturbing number of nightly Internet attacks. At this point I had two goals:

I wanted to thoroughly understand this technology for its own sake. After all, I am fascinated by the issues surrounding Internet security and privacy. I also wanted to understand how everything worked so that I could, if possible, defend against any future similar attacks. (For example, by infiltrating, commandeering, and neutralizing any attacking force of Zombies.)

I also still needed to "get through to" Wicked somehow. I needed to convince him to

lay off his repeated attacks.

My IRC Chat with the ^b0ss^

I had learned a great deal about the Zombies, and I knew that "Wicked" had not created his own as he had claimed. By analyzing the binaries of all the various Zombies my spy-bots had collected, I could pretty much follow the evolutionary "lineage" of this strain of Zombie. I finally found the hacker ("^b0ss^") whose Zombies "Wicked" had "hex edited" in order to create those that had been attacking grc.com.

One afternoon, one of my spy-bots intercepted a conversation taking place between that hacker ("^b0ss^") and another nicknamed "lithium_". Their dialog revealed that "^b0ss^" was creating a new Zombie for "lithium_", editing it to report to a different secret IRC channel using a different password. Unaware that they were under surveillance, they spoke openly of their plans. I didn't discover that interchange until later that evening, but my URL interceptor and downloader had automatically snagged a copy of the new Zombie (this time named "win.exe") and had downloaded it into my Zombie-repository for safe keeping.

Peeking into this new Zombie's now-quite-familiar guts, I immediately noticed something odd: "^b0ss^" had apparently made a small mistake with his Zombie hex editing. He had separated the new strings for the channel and the password key with a period (.) rather than a null (0). This Zombie would not hunt.

I saw an opportunity to help.

Now, don't get me wrong here:
I'm no Zombie lover.

I have NO desire to aid and abet the hackers in their pursuit of unlawful Internet terrorism. "^b0ss^" and "lithium_" would have soon figured out that something was wrong with the new Zombie — and fixed their simple mistake. But I felt that my pointing it out, bringing an olive branch with me into "Zombie-land", would be the perfect means of establishing a constructive dialog with "^b0ss^ (who was likely to be at least somewhat freaked out when I just popped into his private inner-Zombie-sanctum.) So, I ask you to PLEASE KEEP IN MIND that I had a deliberate need and intent underlying the following dialog which took place with "^b0ss^".

When he left one of the main Sub7Server areas and appeared over in the room where he tends his Zombies, I had all of the secret channel names and pass-keys at my disposal. So I just jumped into the room and said "heh". While I was using a Windows IRC chat client, one of my Spy-Bots was automatically recording and logging our dialog from that "long under surveillance" Zombie haven:

```
<Gibson> heh
<^b0ss^> who are you
<Gibson> Hi B0ss.  I'm steve gibson (grc.com) ... ShieldsUP,
<Gibson> OptOut, Leaktest ... and all that stuff.
<^b0ss^> how did you get in here?
<^b0ss^> your not a IRCop
<Gibson> As you might know, my site was attacked (but I don't
<Gibson> think by your bots) a few weeks ago.
<Gibson> Some guy, calling himself "Wicked"
<^b0ss^> my bots?
<^b0ss^> no no
<^b0ss^> I know wicked
<^b0ss^> it was not my bots I promise
<^b0ss^> Wicked has his own
<Gibson> Hey, it's okay
<^b0ss^> alot of bots
<^b0ss^> heh
<Gibson> I know.
<^b0ss^> yeah
<^b0ss^> I promise it wasn't mine
<Gibson> I wanted to let you know that the bot
<Gibson> you made earlier for Lithium would not work
<^b0ss^> what about the bot?
<^b0ss^> you know Lithium
<Gibson> since it has "periods" (2E) instead of NULLS (0)
<^b0ss^> ?
<Gibson> separating the "Channel" and "Key" strings
<^b0ss^> you his friend
<Gibson> no.
<Gibson> I wanted to learn about this stuff
<Gibson> since Wicked was attacking me
<^b0ss^> then how did you know about this place
<^b0ss^> you have your own server?
<Gibson> So I wrote some fake bots to monitor various Bot
<Gibson> networks so that I could learn.
<^b0ss^> damnit
<^b0ss^> so you been spying?
<^b0ss^> hehe
<Gibson> Yeah
<Gibson> But not to worry, I'm no narc.
<Gibson> I just want to be left alone.
<^b0ss^> but how did you get the Key
<^b0ss^> I don't even know you
<^b0ss^> I don't bother anyone with my bots
<Gibson> Check out GRC.COM.  That's me.
<^b0ss^> okay
<^b0ss^> you don't like wicked?
<Gibson> Well ...... I can't say that I know him,
```

```
<Gibson> but he spent a few weeks blasting my site
<^b0ss^> damn
<Gibson> since he thought (he sez that Hellfirez and DrGreen told
<Gibson> him) that I was referring to them as "script kiddies" ...
<^b0ss^> hehe, I got enough bots to blast away a site
<^b0ss^> but I don't use them for that
<^b0ss^> lol
<Gibson> (You have 241 Bots.)
<^b0ss^> thats not it
<^b0ss^> not just on this server
<^b0ss^> how in the hell do you know how many bots I have?
<^b0ss^> damn
<Gibson> I've tracked 241 coming and going over the past four days.
<^b0ss^> let me get some of your bots
<^b0ss^> lol
<^b0ss^> I can't believe this shit, what kind of bot you have
<Gibson> Do you know where Wicked got his?  He claims that he
<Gibson> wrote it, but it looks like a pure hex-edit to me.
<^b0ss^> oh no
<^b0ss^> lol
<^b0ss^> he didn't make them
<^b0ss^> he got his bot from these bots in this room
<Gibson> You really ought to check out my site.  grc.com
<^b0ss^> I am right now
<^b0ss^> ;)
<^b0ss^> nice page
<Gibson> Yeah, I believe that about Wicked.
<Gibson> His channel is #pines1 and PassKey is "penile"
<Gibson> (pines1 is "penis1" with the vowels swapped).
<^b0ss^> lol
<^b0ss^> damn
<^b0ss^> you are pretty good
<Gibson> Anyway, last week I learned IRC protocol and wrote a
<Gibson> bunch of infiltration bots in order to figure out where
<Gibson> all these attacks were coming from.
<^b0ss^> hmmm
<Gibson> It looks like he's lost his dynDNS
<^b0ss^> you know what serve he keeps them all on
<^b0ss^> yup
<Gibson> yeah, I have his server, but I think he's off the air
<Gibson> for now and won't be bothering me again any time soon.
<^b0ss^> we had alot of bots on ips.mine.nu
<^b0ss^> but they took it down
<^b0ss^> for illegal use
<Gibson> cool! I was hoping that might be it.
<^b0ss^> oh, I wouldn't say that
<^b0ss^> he is gettin army back
<^b0ss^> heh
<^b0ss^> I know he has more
```

```
<^b0ss^> somewhere
<Gibson> I don't care if he wants to blast IRC folks,
<Gibson> but I haven't done anything to bother him.
<Gibson> If he blasts me again I'll take them away.
<^b0ss^> lol, he is 13
<^b0ss^> did you know that
<Gibson> Yeah, he said, and he writes like he is.  But I didn't
<Gibson> think he could really write that Bot from scratch.
<^b0ss^> which bot you talkin about
<^b0ss^> do you know mimic?
<Gibson> You call yours "evilbot" (version 0.4c) ...
<Gibson> he renamed it "WkD Bot" (version 1.0)
<^b0ss^> yeah
<Gibson> I don't know anyone. YOU are the first person I've
<Gibson> ever talked to on IRC. Wicked and I have eMailed.
<^b0ss^> mimic has a hell of a bot
<^b0ss^> so, you set up a bot in this channel spying?
<Gibson> Yep .... about a week ago.  I have a list of all
<Gibson> the attacks you've made, etc. etc.
<^b0ss^> shit
<Gibson> The one on a machine within IBM freaked me out.
<^b0ss^> so how did you get the key to my channel to get the bot in
<^b0ss^> IRCop?
<Gibson> Like I said, I just needed to learn about this
<Gibson> stuff so that I could defend myself.
<^b0ss^> man, I wouldn't attack you I promise you that
<^b0ss^> I have no reason
<Gibson> I asked all of the ISP's of the people whose
<Gibson> machines were attacking me for a Bot.
<^b0ss^> oh
<Gibson> Someone sent me one ... and from there I knew what I needed.
<^b0ss^> hehe
<Gibson> Then I wrote a custom "spy bot" and started monitoring
<Gibson> more and more conversations, following leads, URL's, etc.
<^b0ss^> hmmm
<Gibson> that's how I know about you making the new custom bot
<Gibson> for lithium this afternoon.
<^b0ss^> damnit
<Gibson> but when I finally looked at it I saw that it wouldn't work,
<Gibson> so I figured I'd introduce myself and let you know.  :)
<^b0ss^> lol
<Gibson> And of course the Bot itself knows how to logon here! <g>
<^b0ss^> yeah
<^b0ss^> good job
<^b0ss^> I must say
<Gibson> Well, it was nice to meet you.
<^b0ss^> nice to meet you to
<^b0ss^> You are pretty good
<Gibson> And, again, that Bot you made for lith earlier won't work.
```

```
<Gibson> so make sure he doesn't deploy it until you fix it for him.
<^b0ss^> may I ask how old you are?
<Gibson> I'm 46. (Been hacking since I was 14.)
<^b0ss^> lol, alright, thanx
<^b0ss^> damn
<^b0ss^> you are good
<^b0ss^> you gonna leave your bot in here?
<Gibson> Nope.  It's done it's job.  I'm working on a new web page
<Gibson> to talk about the Wicked attacks, and to explain this whole
<Gibson> bizarre world.
<^b0ss^> alright thanx
<^b0ss^> hehe, yeah
<Gibson> Check back at grc.com in a few daze
<^b0ss^> okay
<^b0ss^> I will
<Gibson> If you see Wicked, tell him we had a nice chat
<Gibson> and ask him to lay off.  I don't want to upset him,
<^b0ss^> okay
<Gibson> but I need to, and will, defend my site.
<Gibson> Thanks!
<Gibson> .
<^b0ss^> hehe, okay
<^b0ss^> welcome
```

As we now know,

A 13 year-old hacker . . .

. . . (living in Kenosha, Wisconsin) who goes by the hacker handle "Wicked", was informed by some senior hackers — among them "HeLLfiReZ" a member of the notorious Sub7 crew — that I had referred to them in an online forum, using the derogatory term "script kiddies". I had not. But these senior hackers were upset over a dispute that had erupted in one of our Internet security newsgroups.

"Wicked's" response was to team up with two other hackers, all of whom tend and manage large fleets of "IRC Attack Bots". They launched a concerted and extended "packet attack" against grc.com. In the slang that I learned while monitoring their many conversations, they "packeted" us. They did this, not using any tool they had written, and not possessing the ability to create such a tool themselves, but using a powerful "IRC Bot" that had been passed around extensively. Neither Wicked nor his friends know who wrote it or even where it came from.

The "Wicked" Attacks

"Wicked" and his IRC Bots communicate by logging onto an IRC server located at the domain "wkdbots.***.**" (I have blanked the upper portion of the domain to allow me to provide all other details.) This domain name is hosted by a dynamic DNS service, allowing Wicked to change the location of the IRC server, as needed, by pointing the "wkdbots" domain at a different IP address. This highlights one of the

several weaknesses of the IRC Bots system: A single discovered Bot reveals the IRC meeting place of the entire Bot fleet. The subsequent loss of access to their shared domain cripples the Bot network by denying its access to its central communications hub.

After "Wicked's" Bots connect to the "wkdbots" server, they join and monitor the secret "#pines1" IRC channel using the password key "penile". (You will note that "pines" is the word "penis" with the vowels transposed.) The IRC Bot which Wicked hex-edited to create his own uses a more straightforward channel and password key.

On the evening of May 4th, 2001, after logging onto the wkdbots IRC server and joining his Bots on the #pines1/penile channel, "Wicked" typed the following two commands which were immediately relayed, courtesy of their shared IRC chat server, to his awaiting fleet of IRC attack Bots:

```
!p4 207.71.92.193
!udp 207.71.92.193 9999999 0
```

Within moments, 474 security-compromised Microsoft Windows PC's, containing remote control attack "Zombies", were attacking the grc.com server.

A small sampling of the "Bot's" command language . . .

- **The leading exclamation point "!"** prefixes all of the commands accepted by the IRC Bot.

- **The !p4 command,** followed by the target IP address (in this case the IP of the grc.com server) executes the following standard Windows "ping" command with the following arguments:

```
ping.exe 207.71.92.193 -l 65500 -n 10000
```

This causes the Windows machine to send ten thousand very large (64 kbyte) "ping" packets to the machine at the specified IP. A bit of math shows that this is 655 megabytes of data. This doesn't generate a high-speed stream because the "ping" command waits for a reply before trying again. But if many machines are all pinging at once, the result is cumulative and can be significant. Since the ping command is being executed by a separate (ping.exe) program, the Bot is then free to wreak more havoc . . .

- **The !udp command,** does far more damage. It specifies the target IP, the number of huge UDP packets to send, and the inter-packet delay (zero in this case). The receipt of the !udp command shown above causes each Bot to send 9,999,999 maximum size UDP packets to the grc.com server as fast as the host machine's outbound bandwidth will allow.

In response to receiving the !udp command, each of the hundreds of attacking

machines replied:

PRIVMSG #pines1 :BoMbInG: 207.71.92.193, PaCkEtS: 9999999, DeLaY: 0

. . . and then began "bombing" grc.com with a blizzard of fragmented UDP and ICMP packets, thus consuming our bandwidth several times over and denying our services to the Internet.

A modest number of cable-modem connected home PC's, pumping out maximum-size UDP packets at their maximum upstream bandwidth, were easily able to flood and completely consume grc.com's bandwidth.

This was no "finesse" attack. There was nothing clever about it. The 13 year-old perpetrator had not created the attack tool, and never could. He barely even knew how it operated. But like someone who is handed a loaded gun, even lacking any understanding of how that gun operates, "Wicked" was able to pull the trigger.

GRC.COM was knocked off the Internet for 17 hours by a classic Distributed Denial of Service (DDoS) attack.

Before I leave the topic of Zombie/Bot commands, I thought you might find the **"!r"** command interesting. **"R"** must be short for "Ready" or "Report" or "Rally" because it immediately causes all currently connected and listening Zombie's to "report in". Here is a snippet that one of my spy-bots picked up when "^b0ss^" gave the "!r" command to his troops:

```
<^b0ss^> !r
<xknb> evilbot 0.4c ready for attack...
<oep> evilbot 0.4c ready for attack...
<kvmadj> evilbot 0.4c ready for attack...
<yqvc> evilbot 0.4c ready for attack...
<pvnlz> evilbot 0.4c ready for attack...
<jizl> evilbot 0.4c ready for attack...
<umc> evilbot 0.4c ready for attack...
<wzdr> evilbot 0.4c ready for attack...
<vqfvmh> evilbot 0.4c ready for attack...
<ossqd> evilbot 0.4c ready for attack...
<gyc> evilbot 0.4c ready for attack...
<lvk> evilbot 0.4c ready for attack...
<myv> evilbot 0.4c ready for attack...
<rozjh> evilbot 0.4c ready for attack...
<usxaw> evilbot 0.4c ready for attack...
<vsma> evilbot 0.4c ready for attack...
<xheweq> evilbot 0.4c ready for attack...
<bgpc> evilbot 0.4c ready for attack...
<mntt> evilbot 0.4c ready for attack...
```

```
<nngp> evilbot 0.4c ready for attack...
<mhonm> evilbot 0.4c ready for attack...
<fgjc> evilbot 0.4c ready for attack...
<gyk> evilbot 0.4c ready for attack...
<mkenx> evilbot 0.4c ready for attack...
<pnyd> evilbot 0.4c ready for attack...
<btkh> evilbot 0.4c ready for attack...
<qwbbd> evilbot 0.4c ready for attack...
<vst> evilbot 0.4c ready for attack...
<griv> evilbot 0.4c ready for attack...
<frf> evilbot 0.4c ready for attack...
<fdhcmmk> evilbot 0.4c ready for attack...
<kdu> evilbot 0.4c ready for attack...
<jsea> evilbot 0.4c ready for attack...
<yxkoo> evilbot 0.4c ready for attack...
<xcchvl> evilbot 0.4c ready for attack...
<elzll> evilbot 0.4c ready for attack...
<imebw> evilbot 0.4c ready for attack...
<jddp> evilbot 0.4c ready for attack...
<bpfzvk> evilbot 0.4c ready for attack...
<egja> evilbot 0.4c ready for attack...
<nqab> evilbot 0.4c ready for attack...
<fzdnt> evilbot 0.4c ready for attack...
<eyddzxs> evilbot 0.4c ready for attack...
<dshraf> evilbot 0.4c ready for attack...
<xigvx> evilbot 0.4c ready for attack...
<iaamew> evilbot 0.4c ready for attack...
<joog> evilbot 0.4c ready for attack...
<fhtveo> evilbot 0.4c ready for attack...
<rfktnd> evilbot 0.4c ready for attack...
<ojjntjw> evilbot 0.4c ready for attack...
<esx> evilbot 0.4c ready for attack...
<gxbgi> evilbot 0.4c ready for attack...
<hhenh> evilbot 0.4c ready for attack...
<wknt> evilbot 0.4c ready for attack...
<hwn> evilbot 0.4c ready for attack...
<nnk> evilbot 0.4c ready for attack...
<brmbpb> evilbot 0.4c ready for attack...
```

Those random-character evilbot nicknames, appearing in angle brackets at the front of each line, are the random names generated by each Zombie when it logs on to the IRC server. In the unlikely event of a nickname collision the Bot simply generates another.

The Birth of the IRC Bot (Windows Attack Zombie)

The computer press calls them "Zombies", but they are known as "Bots" or "IRC Bots"

within the hacker community. The particular strain of IRC Bot that attacked us calls itself an "evilbot".

"IRC Bots" are among the newer breed of Distributed Denial of Service (DDoS) agents deployed by the Internet's most active hackers. Whenever an IRC Bot hosting Windows PC is started, the Bot waits for the system to finish booting, then connects to a previously designated IRC server. Using a private password key, it joins a secret IRC channel that is not visible to other users of the IRC server . . . and awaits commands.

The use of a central IRC server provides significant anonymity, deployment, and logistical benefits to the hackers, although as we have seen, it is not without liabilities. I was able to successfully penetrate the Bot's world without much trouble.

IRC servers are often configured to deliberately obscure the names and IP addresses of their clients, thus providing anonymity to all users. Since this anonymity can only be breached through physical access to the server, many Bot armies are "run" from servers located on foreign soil where access is impossible to obtain.

Since IRC Bots "phone home" to the central server, the hacker does not need to know which specific machines are hosting his Bots. This allows Bots to be deployed in the wild by a wide variety of means. Hackers create Bot-carrying eMail viruses (frequently enabled by Microsoft's virus-friendly Outlook Express), they create infected Internet "Trojan" downloads, place Bots in USENET newsgroups, and do anything they can to get their Bots into other people's computers.

IRC Bots never need to be "scanned for" since all active Bots contact their home base IRC server whenever they "awaken". The various IRC Bots I have acquired and examined are just 15,904 bytes in size, so they are easily hidden as trojans within other, typically huge, Windows programs.

A Quick & Easy Check for IRC Zombie/Bots

If you have managed to read all the way through this lengthy and detailed adventure, I am sure you will agree that you do NOT want any of these nasty Zombies or their relatives running around loose inside your PC. Fortunately, it's quite easy to verify that your system is not currently infected by one of these IRC Zombie/Bots.

All of the IRC Zombie/Bots open and maintain static connections to remote IRC chat servers whenever the host PC is connected to the Internet. Although it is possible for an IRC chat server to be configured to run on a port other than "6667", every instance I have seen has used the IRC default port of "6667".

Consequently, an active connection to an IRC server can be detected with the following command:

```
netstat -an | find ":6667"
```

Open an MS-DOS Prompt window and type the command line above, then press the "Enter" key. If a line resembling the one shown below is NOT displayed, your computer does not have an open connection to an IRC server running on the standard IRC port. If, however, you see something like this:

```
TCP   192.168.1.101:1026   70.13.215.89:6667  ESTABLISHED
```

. . . then the only question remaining is how quickly you can disconnect your PC from the Internet!

A second and equally useful test can also be performed. Since IRC servers generally require the presence of an "Ident" server on the client machine, IRC clients almost always include a local "Ident server" to keep the remote IRC server happy. **Every one of the Zombie/Bots I have examined does this.** Therefore, the detection of an Ident server running in your machine would be another good cause for alarm. To quickly check for an Ident server, type the following command at an MS-DOS Prompt:

```
netstat -an | find ":113 "
```

As before, a blank line indicates that there is no Ident server running on the default Ident port of "113". (Note the "space" after the 113 and before the closing double-quote.) If, however, you see something like this:

```
TCP     0.0.0.0:113     0.0.0.0:0     LISTENING
```

. . . then it's probably time to pull the plug on your cable-modem!

> **Note** that a Windows IRC client program running in the PC **will** generate false-positive reports since these are tests for IRC client programs. So be sure to completely exit from any known IRC client programs BEFORE performing the tests above.

Personal Firewalls and IRC Zombie/Bot Intrusions

- **ZoneAlarm v2.6** (Free) —

The last of my testing was to see whether the firewall I keep telling everyone to use: **ZoneAlarm** — either FREE or Pro — would be effective in stopping the IRC Zombie/Bot and the Sub7 Servers that had taken up residence in my poor "Sitting Duck" laptop.

I downloaded the current, completely free, version of ZoneAlarm 2.6 from the ZoneLabs web site and installed it on the "Sitting Duck" laptop. Upon restarting the machine I was gratified to receive immediate notification that the Zombie/Bot was attempting to make an outbound connection to its IRC chat server.

Meanwhile, the Sub7 Trojan was sitting quietly waiting for someone to connect to it. So I used another machine to "Telnet" to the port the Sub7Server Trojan was listening on. Up popped ZoneAlarm asking whether the nonsense-looking random character name the Sub7Server had chosen for itself should be allowed to accept a connection from the Internet.

Perfect performance from ZoneAlarm.

Then I had a thought: What would Network ICE's BlackICE Defender do under the same circumstances?

BlackICE Defender v2.5 ($39.95) —

I did not have a current copy of BlackICE Defender around, but I felt that this was an important test. So I laid out $39.95 through Network ICE's connection to the Digital River eCommerce retailer and purchased the latest version (v2.5) of BlackICE Defender hot off the Internet. I had already removed all traces of ZoneAlarm and restarted the machine, so I installed BlackICE Defender, let everything settle down, and restarted the machine with my packet sniffer running on an adjacent PC.

As far as I could tell, BlackICE Defender had **ABSOLUTELY NO EFFECT WHATSOEVER** on the dialogs being held by the Zombies and Trojans running inside the poor "Sitting Duck" laptop. I knew that BlackICE Defender was a lame personal firewall, but this even surprised me.

The Zombie/Bot happily connected without a hitch to its IRC chat server to await further instructions. The Sub7 Trojan sent off its eMail containing the machine's IP and the port where it was listening. Then it connected and logged itself into the Sub7 IRC server, repeating the disclosure of the machine's IP address and awaiting port number. No alerts were raised, nothing was flashing in the system tray. The Trojans were not hampered and I received no indication that anything wrong or dangerous was going on.

I took a lot of grief after my LeakTest utility cut right through BlackICE Defender. Network ICE told everyone that LeakTest was "being allowed through" because it was a completely benign Trojan. I knew that was a load of bull (and they must have too), but it didn't really matter to me, and I had no affirmative means of proving otherwise.

Well . . . I have that now, and so do you.

I performed one final test: As I had with ZoneAlarm, I attempted to connect to the Sub7Server Trojan running inside the "Sitting Duck" machine on the IP and listening port number the Trojan was advertising all over the Internet **. . .** and it worked perfectly. I received Sub7's "PWD" prompt asking me to login.

Anyone want an "only used once" copy of BlackICE Defender?

I certainly have no use for it.

> **To anyone who is still stubborn enough to insist that BlackICE Defender is actually good for something:** PLEASE do not write to me. I don't want to hear it. I'm a scientist who will not find your mystic beliefs to be compelling. I respect your right to your own opinions, no matter how blatantly they fly in the face of logic and reality. That is, after all, the nature of faith. Happy computing. I suggest prayer.

What's Next?

I have concluded this research into the motivations behind, and the technologies of, the multiple Distributed Denial of Service (DDoS) attacks that were launched against GRC.COM during the month of May, 2001.

- I believe that I have learned everything there is to learn from these IRC Zombie/Bot style attacks. I have found and conversed with the important players, and I have analyzed their tools, technologies, and networks. I could spend the rest of my life pursuing the specifics of all possible exploits and toolz, but that's not where I want to go.

- I have learned that our industry's leading consumer ISP's are worse than useless when asked for any form of help relating to Internet security or the welfare of the paying customers. For reasons unfathomable to me they choose to disavow responsibility for the conduct of their users, and equally refuse to offer any help for their customers' Trojan-infected machines.

- I learned some bitter truths and realities about the nature of Federal government involvement. There are just too many large problems for the smaller ones — which may be destined to grow larger — to receive help or attention. I sincerely hope that the 13 year-old hacker known as "Wicked" figures out that his youthful shield will dissolve in five years. I believe that in the future the United States government, and the world at large, is going to become increasingly intolerant of Internet hacking. The penalties for transgression will be onerous.

- And finally, this experience has reaffirmed my commitment to two separate but closely related needs:

  - **The threat represented by Microsoft's forthcoming Windows XP operating system**, with its confirmed ability to easily generate malicious Internet traffic — for NO good reason — can not be overstated. The proper executives within Microsoft MUST be reached with this message so that those plans can be reviewed in light of the potential for their system's massive abuse of the inherently trusting Internet.

    For more information about this serious threat to the Internet, please see my **WinXP & DoS page**.

  - **The days of an Internet based upon mutual trust among interconnected networks has passed.** The Internet's fundamental infrastructure MUST BE SECURED before the Net becomes further threatened by increasing levels of malicious attacks. Since this requires irresponsible ISP's — who repeatedly demonstrate that they could not care less — to assume the sorts of local responsibility that they have so far deliberately shunned . . .

    We need a tool to hold ISP's accountable and publicly demonstrate individual ISP irresponsibility.

    Given the universal reluctance they have demonstrated so far, I believe that only active public scrutiny will bring about the changes required to insure a reliable and secure future for the Internet.

The development of that tool is my next project.

The name of that FREE tool will be:

Spoofarino tm

(If you would like to be notified when Spoofarino is ready, you are invited to join our user-managed eMail system. You can leave any time)

Purchasing Info | GRC Mail System | To GRC's Home | Tech Support | Discussions

The contents of this page are Copyright (c) 2001 by Gibson Research Corporation. SpinRite, ChromaZone, ShieldsUP, NanoProbe, the character 'Moe' (shown above), and the slogan "It's MY Computer" are registered trademarks of Gibson Research Corporation (GRC), Laguna Hills, CA, USA. GRC's web and customer privacy policy.

~ ~ ~