

Re: Wish list

Source: <http://www.derkeiler.com/Newsgroups/comp.os.linux.security/2005-12/msg00050.html>

- *From:* ibuprofin@xxxxxxxxxxxxxxxxxxxxxxxxxxxx (Moe Trin)
 - *Date:* Thu, 01 Dec 2005 14:50:09 -0600
-

On Thu, 01 Dec 2005, in the Usenet newsgroup comp.os.linux.security, in article <f9idnSSPNdWtzRLeRVn-sQ@xxxxxxxxxx>, Newsbox wrote:

>You are undoubtedly 100% correct. I was looking for a fairly easy, quick
>and accurate way to specific information. Your generalization is
>at least fast and correct, I believe.

Unfortunately. The stuff tends to morph fairly quickly, so many lists that are available tend to be somewhat behind.

>Thanks for the tip. I was actually already familiar with "whois", but
>possibly not all readers were. And then there's the part about the low
>contact success. It's sad, but "whois" really isn't much help in most
>cases, as necessary a utility as it may be. You have given better tips in
>the past.

Simple explanation – whois gets you to the entity assigned the address space. If the address is a major ISP such as 'comcast' or 'telus' or similar, it is possible (but unlikely) that a message will be sent to the box owner. You won't be informed, due to privacy laws and the like. The EU is fairly strict about this, and some European ISPs have taken that as a reason to /dev/null complaints about users. I used to get decent results to complaints sent there. Not any more. If the address is an Asian ISP – Korea is particularly bad – even if someone reads the complaint, it will be tossed.

>Right! Agreed. I guess I let myself open for that blast.

Sorry – the thing that triggered that response was the indication you wanted to counter-attack. "how to exploit any such known vulnerability" is just plain bad procedure.

>Please kindly read my response to Greg Metcalfe's kind (first) message,
>(written yesterday but only posted today)

and after I had written my response and put it into the queue. That was also after your response to matt_left_coast, and Greg and Dale's reply.

>I didn't do the research you ordered. "Fill out a form." I am already

Re: Wish list

> somewhat familiar with the subject matter to which you refer. Your
> presumption is (I take it) intentionally offensive. If your intention was
> not to offend, you failed and I apologize for a curt (if correct)
> answer. If you seriously want to talk about spoofing, start a new thread.

See the response to Greg Metcalfe above. Most UDP spam seems to be from spoofed addresses or zombies where the owner is clueless, the ISP ignoring complaints, or possibly from rentals where the provider is protected as long as they are not pissing off the national authorities.

The TCP stuff is much more difficult to spoof because of the 3-way handshake and sequence numbers, and in my experience is a worm (or skript kiddies acting like one). The two defenses here are not having un-needed ports open, and for those ports that are, having servers up to date and configured properly.

>> Your firewall is blocking this shit – IGNORE IT. You are not the
>> mighty avenger who is going to clean up the world.
>
> With all due and sincere respects and no hard feelings (plus continued
> appreciation of much good help in the past), this does not show your best
> qualities.

That reaction was about your counter attack concept. In case you haven't twigged, that line really set me off, primarily because I see this all too often and like you, I have better things to do. The advice remains. There really isn't that much you can do to counter the attacks and spam coming in from the world. You can't have anyone arrested in country \$FOO where the packets came from, because the "attacker/spammer" isn't likely to be there, never mind that the authorities tend to put this form of crime rather low on the priority list. I'd love to see it happen, but it's long past the time when that is possible. Regarding the messenger spam, the problem a normal user has is that no matter what their firewall does, it's wasted bandwidth. The spammer isn't running a normal daemon, and so sending a RST or ICMP Type 3 (even assuming the address isn't spoofed) isn't going to do anything productive – the spammer's message has been delivered, whether you saw it or not. At work, we've gone to the extent of portshifting our outbound DNS queries out of the range 1025 to (say) 1050, so there won't be any legitimate inbound traffic on those ports – then made arrangements with our upstream to simply drop all inbound UDP in that port range. This saves on our bandwidth, though we've got to pay the provider for the service. (Like many, we are charged on a traffic basis.)

Old guy
.

-
- *Follow-Ups:*
 - ◆ *Re: Wish list*

Re: Wish list

◇ *From:* Newsbox

• **References:**

◆ **Re: Wish list**

◇ *From:* Moe Trin

◆ **Re: Wish list**

◇ *From:* Newsbox

- Prev by Date: **Re: md5 collision**
- Next by Date: **Re: md5 collision**
- Previous by thread: **Re: Wish list**
- Next by thread: **Re: Wish list**
- Index(es):
 - ◆ **Date**
 - ◆ **Thread**