

Re: I have been hacked (WAS: Have I been hacked or is nmap wrong?)

## Re: I have been hacked (WAS: Have I been hacked or is nmap wrong?)

---

*Source:* <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/questions/2006-01/msg01629.html>

---

- *From:* Will Maier <[willmaier@xxxxxxx](mailto:willmaier@xxxxxxx)>
  - *Date:* Wed, 18 Jan 2006 14:58:14 -0600
- 

On Wed, Jan 18, 2006 at 05:38:50PM +0200, Kilian Hagemann wrote:

> On Wednesday 18 January 2006 16:25, Will Maier pondered:  
>> On Wed, Jan 18, 2006 at 03:56:32PM +0200, Kilian Hagemann wrote:  
>>> I have never even heard of "frox" before, but after some  
>>> googling it turns out that it's a GPL'ed transparent ftp  
>>> proxy...  
>>  
>> Where's it pointing?  
>  
> No idea, I only went as far as trying to login anonymously using a  
> console based ftp client. How could I find out?

Connect to it and watch the packets in tcpdump(8) or similar. this may not give you the full answer, but it'll help. What banners do the FTP servers have? Is there a domain listed? Who owns that domain?

>> What do you see when you connect to the SMTP ports? Are they  
>> really mail servers, or just rogue services running on 25?  
>  
> They are really mail servers, at least smtp for outgoing mails  
> (don't know about incoming though). I used kmail to configure them  
> as standard outgoing smtp mail servers and successfully sent  
> myself two emails, one via each server. Surely a default, out of  
> the box, unconfigured and sendmail\_enable="None" sendmail process  
> wouldn't allow for something like that, never mind the fact that  
> the firewall is supposed to block ANY access from the outside  
> (output of ipfw show is attached)

So these are running, functioning sendmail servers that /you/ didn't configure (on purpose)? What do you see when you 'talk' to them via nc(1)? If you're firewall was dropping incoming packets destined to those ports, you wouldn't have been able to send a mail through them (or connect on 25 with nc(1))...

> Well, I didn't worry about samba because it's firewalled to the  
> outside(unless some Windows virus on one of the LAN machines  
> exploited a samba hole, is that likely?).

Re: I have been hacked (WAS: Have I been hacked or is nmap wrong?)

Re: I have been hacked (WAS: Have I been hacked or is nmap wrong?)

I don't know Samba that well, but it's possible it could be exploited (check the web for recent advisories pertaining to it).

How much do you trust the users on the 'green' side? Could one of their boxes have been compromised and then used as a platform to attack your border servers? This sort of (nightmare) scenario is why people have been whining about 'defense in depth' for the last few years; it turns out that your crunchy, impermeable outside actually can be as squishy as your inside.

- > There is only one single normal user account with an uncommon name
- > and an impossible password(16 characters randomly generated from
- > ASCII charset). ChallengeResponseAuthentication is commented out
- > in sshd which I guess means it uses the standard PAM
- > authentication. It also allows password/interactive authentication
- > in addition to public key, I always use the former. I do admit
- > that I have set "PermitRootLogin yes" but my root password is 9
- > characters with numbers and non-alphanumeric characters, so hard
- > to brute-force.

Having a kickass, long username with an 'impossible' 16 char password and an open root account with a password 9 chars long is like putting a heavy steel door on a cardboard box. Allowing PermitRootLogin is a mistake in almost every scenario; disable it in the next generation of your servers (if possible). It's a 'weakest link' sort of situation, I guess.

- > In any case, it's important to note that the only access from the
- > outside via ssh/rsync is firewalled in such a way that it only
- > allows access from a single IP address which my institution
- > assigns me statically via DHCP (see attachment).

That's good.

- > They would have had to a) find out what this one and only trusted
- > IP address is b) spoof it successfully c) attack ssh brute force?

Assuming the firewall works, they would certainly have to complete steps a, b and c; unless, that is, they compromised /your/ box, too. Unlikely, though, I suspect.

- > Well, I thought my setup was secure enough for a very basic
- > router/gateway/firewall for a couple of Windows machines using a
- > sucky internet connection which is not worth stealing.

Unfortunately, the asset you should be protecting might not be your bandwidth or data or whatever it is you've been assuming. When you set up a firewall, you're protecting something — in your case, what is it? Have you defined that for yourself? It's hard to do a good job defending something you haven't or can't define. While it probably sounds pedantic or silly, take a moment to ask yourself

Re: I have been hacked (WAS: Have I been hacked or is nmap wrong?)

Re: I have been hacked (WAS: Have I been hacked or is nmap wrong?)

what it is you want to protect. If there are several things, rank them by priority. \_Then\_ go about designing a defense. Securing your stuff may not be a terribly high priority at all; if so, accept the fact that something bad will happen once in a while. Your security plan might just be "deal with it when the shit hits the fan." No problem. That can make sense. Having clarified that for yourself, though, makes things easier.

> So I didn't go through the effort of using a file integrity  
> monitor, remote logging, traffic dumps or network monitors (jeez,  
> sysadmins lives are really difficult these days :-())

Like I said above, those sorts of defenses might be overkill for you. That's fine — there are benefits, though, to deploying them, and their cost might not seem so bad in context of your (evident) frustration at the break-in. I still sympathize, by the way: it's a tricky, awful world sometimes.

> Thanks a ton for the help and advice, I'll see what I can do.

Of course — best of luck with damage control.

> # output of ipfw show. I edited it to remove all count rules(merely used for traffic accounting),  
> # some unreachable rules to prevent some LAN clients from accessing the internet altogether  
> # and substituted some ip blocks for privacy purposes. LAN\_NET is the LAN subnet,  
> # MY\_OUTSIDE\_IP is the unique and only ip address that is allowed to login from the outside via  
ssh/rsync

I don't use IPFW, so I can't really help here. You might want to repost your config to a more firewall-oriented list if you don't get enough response here; they'll certainly be able to help you.

--

o-----{ Will Maier }-----o  
| jabber:..wcmaier@xxxxxxxxxxxxx | email:.....wcmaier@xxxxxxx |  
| \.....wcmaier@xxxxxxxxxxxxx | \.....wcmaier@xxxxxxxxxxxxx |  
\*-----[ BSD Unix: Live Free or Die ]-----\*

freebsd-questions@xxxxxxxxxxxxx mailing list

<http://lists.freebsd.org/mailman/listinfo/freebsd-questions>

To unsubscribe, send any mail to "freebsd-questions-unsubscribe@xxxxxxxxxxxxx"

---

• **References:**

- ◆ ***Have I been hacked or is nmap wrong?***  
    ◇ From: Kilian Hagemann
- ◆ ***I have been hacked (WAS: Have I been hacked or is nmap wrong?)***  
    ◇ From: Kilian Hagemann
- ◆ ***Re: I have been hacked (WAS: Have I been hacked or is nmap wrong?)***

Re: I have been hacked (WAS: Have I been hacked or is nmap wrong?)

Re: I have been hacked (WAS: Have I been hacked or is nmap wrong?)

◇ *From:* Will Maier

◆ ***Re: I have been hacked (WAS: Have I been hacked or is nmap wrong?)***

◇ *From:* Kilian Hagemann

- Prev by Date: ***Why I haven't device /dev/cd0***
- Next by Date: ***Re: Why I haven't device /dev/cd0***
- Previous by thread: ***Re: I have been hacked (WAS: Have I been hacked or is nmap wrong?)***
- Next by thread: ***Re: I have been hacked (WAS: Have I been hacked or is nmap wrong?)***
- Index(es):
  - ◆ ***Date***
  - ◆ ***Thread***