

Re: tcp vulnerability? haven't seen anything on it here...

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2004-04/4827.html>

From: Willy Tarreau (w_at_w.ods.org)

Date: 04/22/04

Date: Thu, 22 Apr 2004 16:18:48 +0200
To: "Richard B. Johnson" <root@chaos.analogic.com>

Richard,

you are confusing several things, stateful vs stateless protocols. A ping doesn't need a session on the remote host to be interpreted. A TCP segment whose flags don't show a SYN need a session to be interpreted. Please note that I'm not arguing that you won't crash a linux box with an RST addressed to a broadcast address, I'm saying that there's absolutely no reason why this should reset all connections, as you proposed it. Someone would have had to code this explicitly, it cannot be a simple side effect.

Imagine that each packet which enters the system is presented to a hash table containing the sessions, and that its session is looked for into this hash table. You agree that in such code, there's no reason to find anything that runs through all sessions and kill everyone, since this code has no use there, and has no reason to be implemented on purpose !

Look at functions such as `tcp_v4_lookup()` in `net/ipv4/tcp_ipv4.c` for example. When it reaches `tcp_v4_lookup_established()`, you find this :

```
for(sk = head->chain; sk; sk = sk->next) {
    if(TCP_IPV4_MATCH(sk, acookie, saddr, daddr, ports, dif))
        goto hit; /* You sunk my battleship! */
}
```

You cannot match more than once.

Cheers,
Willy

On Thu, Apr 22, 2004 at 09:42:58AM -0400, Richard B. Johnson wrote:

> On Thu, 22 Apr 2004, Willy Tarreau wrote:

>

> > On Thu, Apr 22, 2004 at 07:35:54AM -0400, Richard B. Johnson wrote:

> >

> > > Has anybody checked to see what Linux does if it receives a

Linux–Kernel: Re: tcp vulnerability? haven't seen anything on it here...

> > > *RST to the broadcast address? It would be a shame if all*
> > > *connections were dropped!*
> >
> > *I don't see how this would be possible : a TCP packet is matched *only* if*
> > *it refers to a valid session. If you have no session established from/to the*
> > *broadcast address, there's no possibility that an RST targetted at this*
> > *address*
> > *terminates anything, even if the ports are OK.*
> >
> > *Cheers,*
> > *Willy*
> >
>
> *If course it's possible. Remember the trick to blue–screen W\$, just*
> *send a fragmented packet with a large length, then never send the*
> *rest. There are lots of things that can happen when control*
> *data goes to the broadcast address. Ping the broadcast address and*
> *observe. If you have any W\$/2000/prof machines on your network that*
> *don't have service–pack 2 or later installed, just syn–flood the*
> *broadcast address. So I wonder how well the corner cases have been*
> *checked. Of course you can't "connect" to a host using the broadcast*
> *address, unless some code runs off the end of a switch statement*
> *unchecked.*
>
> *Hopefully invalid packets just get dropped on the floor. However,*
> *history shows otherwise. Linux has a habit of loudly complaining*
> *about invalid packets or protocol violations. The result being*
> *a log full of messages leading to a full file–system. Fortunately*
> *one can turn off many using the /proc/sys/net/ipv4 interface.*
>
> *Cheers,*
> *Dick Johnson*
> *Penguin : Linux version 2.4.26 on an i686 machine (5557.45 BogoMips).*
> *Note 96.31% of all statistics are fiction.*
>
–

To unsubscribe from this list: send the line "unsubscribe linux–kernel" in
the body of a message to majordomo@vger.kernel.org
More majordomo info at <http://vger.kernel.org/majordomo–info.html>
Please read the FAQ at <http://www.tux.org/lkml/>