

Re: Changes in IDS Companies?

Source: <http://www.derkeiler.com/Mailing-Lists/securityfocus/focus-ids/2002-10/0143.html>

From: Matt Harris (mdh@unix.si.edu)

Date: 10/31/02

Date: Thu, 31 Oct 2002 10:48:54 -0500

From: Matt Harris <mdh@unix.si.edu>

To: Aaron Turner <aturner@pobox.com>, focus-ids@securityfocus.com

Aaron Turner wrote:

>

> On Tue, Oct 29, 2002 at 09:28:08AM -0500, Matt Harris wrote:

>>

>>

>> Aaron Turner wrote:

>

>>> 1) Futzng with router ACL's or firewall policies via your IDS is not granular.

>>> They don't drop a specific connection (the attack) but rather all traffic on

>>> a given port for a client/server. This can have very ugly effects for

>>> legit traffic.

>>

>> Generally, this is done on a basis of simply blocking all inbound

>> traffic from the offender's IP address. Hence entirely blocking the

>> effective attack as well as anything else they may try for the next X

>> number of seconds/minutes/whatever.

>

> That's exactly what you shouldn't be doing. Let's say you detect someone

> attacking your network. How do you know:

>

> 1) That the packets don't have a forged source IP?

Is there a reliable way to discover/deal with this? I haven't really seen/read much on this subject, though it sounds like, from what I have read, it's so hard to do that very few people have the facilities to do it at their disposal.

>

> 2) The the user isn't behind some HTTP, socks, etc proxy?

Good. I want to block misconfigured proxies. Who wouldn't? :-)

>

> Either case and you've likely killed perfectly legit traffic while
> stopping the attack, perhaps preventing paying customers from doing
> business with you. Things like port scans and DoS attacks very often

Re: Changes in IDS Companies?

SecurityFocus IDS: Re: Changes in IDS Companies?

> *include packets with forged source IP's.*

We have no paying customers – we're a research institution within the federal government. As far as the attacks possibly having forged sources, again, what is a good way to deal with this potential? I don't see that an inline NIDS would really be able to do any more about these than a non-inline NIDS. Correct me if I'm wrong – I'm definitely no NIDS guru. :-)

>

> > > 2) *It's too late. The attack has already reached the target. Consider something like jill.c which exploits the IIS-ISAPI buffer overflow and opens a connection back to the attacker on another port and you'll quickly understand why this method of "protection" is more hype than reality.*

> >

> > *If people are running insecure web servers, then is it really the network infrastructure's job to protect them?*

>

> *I've never met any admin of any OS (Solaris, Linux, Windows mostly) who claimed that he/she had patched all of the servers within 24 hours of a patch on a regular basis. Most wouldn't even claim 7 days or even a few weeks. Is this best-practices? Not even close. Is it the reality? Absolutely, especially since most companies don't have their IT group fully staffed due to the economy.*

We're not an IT company. :-) I'm not saying that the solution that I'm designing is right for everyone, only that it has done very well thus far.

>

> *When you consider most (all??) worms effecting IIS were exploiting bugs which had patches released months in advance, it's clear to me at least that companies are either unwilling or unable to keep up. Hence, it seems reasonable that the market will come up with an alternative solution which requires less effort on the admin. (Assuming they don't all move their servers to OpenBSD :-)*

Pretty sad state of affairs, when people don't update their patches at least once a month. I do. If other people don't, it seems to me that they possibly or probably will get broken into. Again, I'm glad that I do. :-)

>

> > *I'm thinking more along the lines of protecting against flood attacks, port scans, and the like – from smurfs to simple icmp floods, etc. In addition, blocking at the border router level can be even more useful for this, since it stops it before it gets to the IDS, Firewall, etc, and hence saves them some processing time for legitimate traffic. It's not a perfect solution to all problems, but IMO the only real solution has to be at every level – I only go so far with network based security, and rely on host based*

SecurityFocus IDS: Re: Changes in IDS Companies?

- > > *security for the rest. Exploits just shouldn't work against systems,*
- > > *and if they do because some admin was lazy, then it shouldn't be my*
- > > *IDS's job to protect their lazy selves. ; -)*
- >
- > *While I want to agree with you (there's something nice in the thought that*
- > *only lazy admins get their servers broken into), in reality it's not a*
- > *question of laziness. Generally I see a few major issues:*
- >
- > *1) Just not enough people to do all the work. The economic downturn makes*
- > *this even worse than it was with many companies laying people off or*
- > *imposing hiring freezes.*

I am inclined to agree here. But at the same time, doesn't this simply make it clear that, at least for the 99% most part, "only lazy people or stupid companies" get broken into?

- >
- > *2) Too many patches and servers to keep up with. Just trying to keep*
- > *up with all the security patches that the vendors keep spewing is insane.*

That's why most vendors nowadays have released automation systems (or at least engineered their patches in such a way that lends itself to automating it). For example, once a month, I hit sunsolve, download a cluster of patches that I want selectively, and since all of my servers are running my standard build (and I won't build anything differently), I push them out via sdist to every server, have them applied, and reboot each server in turn based upon the schedule I've set out to everyone who'll be kicked off their applications by me performing those reboots. Sun also provides their own automation methods, but in my case, I created this system before theirs' was mature, and I like the way my system works. Doing it once a month has never been a bad thing, and I keep apprised of emergency patches and such which I can apply separately as needed.

- >
- > *3) Also, some very popular vendors *cough*Microsoft*cough* like to*
- > *downplay the vulnerability to save face, so admins even if they are trying*
- > *to keep up tend to prioritize patches poorly.*

This is a vendor problem, then, and people should stop using vendors who do not meet their security needs. :-)

- >
- > *4) Patching systems often cause downtime. Hence, it often requires the*
- > *work to be done during non-peak hours (late at night). IT people,*
- > *contrary to popular belief do occasionally have a life/family and can't be*
- > *doing patches 7 nights a week (assuming their windows would even allow that).*

I usually push out my patches during the day, then cycle through and reboot the critical systems between 5 and 6 pm (I work 9-6 anyways, so it doesn't inroach outside of my schedule), and then reboot the

SecurityFocus IDS: Re: Changes in IDS Companies?

non-critical systems that no one will notice anyways throughout the next day. A reboot of a system shouldn't take more than an hour, and none of mine take more than 15 minutes at most on the really big ones with tons of disk arrays. :-)

>
> 5) *Plain ignorance and/or laziness. Yes, some admins think it'll never happen to them and that nobody would ever target them. We all know they're wrong, and get pissed off when it's now their servers attacking us.*

And we should have firewalls to block out their servers. If they're internal, we should have our network folks turn off their switch ports. :-)

>
> > *Security is everyone's concern. If it isn't a particular person's concern, then they'll be the ones to have to fix or rebuild their systems.*
>
> *Yep. Of course as many people have been arguing, security should be done in depth. I'm not saying an NIPS can prevent all attacks so you don't have to ever patch your systems again. That's insane.*
>
> *I tend to think of inline NIPS as a lifejacket. If you're smart and pay attention, you really shouldn't ever need it. But if something bad happens, it's a real good thing to have. And of course, if you're really stupid or just unlucky, even a lifejacket won't save you.*

That's exactly how I think of my IDS systems – just a different physical architecture is all.

>
> > *But that's a philosophical and business difference for a lot of people. I'm in a place where business decisions don't affect things since we're not running a business. And as far as philosophy, see above.*
>
> *Consider yourself lucky then! Not many of us can say that business decisions don't effect our work.*

Non-technical [read: business] people have no place making decisions that affect technical systems. Technical people will provide a very specific, and very locked down service to the business people (ie a web server with ports 22, 80, and 443 open to it, and which will be patched on our regular rotation). They can use SCP to upload their content as an unprivileged user. Seems like a relatively secure configuration to me, as long as it isn't running IIS, of course. And why would they need anything else? The business types like to make up BS justifications for insecure applications, and they really rely on the technical community to smack them down when they ask/tell us to allow that sort of thing to go on. But I haven't worked in a commercial environment for close to two years now, maybe it's changed. I've found that BOFH style network

SecurityFocus IDS: Re: Changes in IDS Companies?

administration always works best and keeps the users happiest, though. Especially when they hear their friends complaining about viruses and exploits and such and don't see any of that themselves. ;-)

>
> > > 3) Many attacks are internal. Most firewalls are at the border, hence
> > > there's nothing the firewall can do, unless you (re)deploy more firewalls.
> >
> > True enough. Deploying internal firewalls and IDS's is definitely not a
> > bad thing, if not in fact even a good thing. Most of the attacks I see
> > internal are unintentional user-mishaps, I've yet to see any genuine
> > malicious activity. But nonetheless, we try to be prepared.
> > Statistically here, about 99% of attacks outside of individual subnets
> > (I have no way of monitoring what may go on within a separate subnet,
> > though I think the help desk would be getting calls if something bad
> > happened that affected users adversely), come from the internet. So,
> > that is where the effort here is in fact concentrated.
>
> Expecting your help desk to notice/get calls is a big if. An obvious
> example was the latest attack on the root name servers. Definitely
> an attack, just most people didn't happen to notice. The root
> name servers are closely monitored by the admins of course, so they
> knew even if the users didn't.

Anything that users notice, they will call in about – they love to complain. :-)

And anything that will affect my UNIX boxes, I will get an automated cellphone page about – my systems love to complain, as well. :-)

>
> Consider the IIS-ISAPI exploit again... since IIS restarts after it crashes
> unless someone was paying attention to the logs (or had an IDS) one would
> generally not realize they had been broken into.

Sounds like an architectural flaw in microsoft's design. Not really something that someone concerned with providing a secure infrastructure should be concerned about, if the end-users are running poorly designed systems/software.

--
/*
*
* Matt Harris - Senior UNIX Systems Engineer
* Smithsonian Institution, OCIO
*
*/

-
- **Previous message:** [Marsu Pilami: "Log correlation"](#)
 - **In reply to:** [Aaron Turner: "Re: Changes in IDS Companies?"](#)
 - **Next in thread:** [J. Foobar: "Re: Changes in IDS Companies?"](#)
 - **Next in thread:** [Marcus J. Ranum: "Re: Changes in IDS Companies?"](#)

SecurityFocus IDS: Re: Changes in IDS Companies?

- *Messages sorted by:* [date] [thread] [subject] [author] [attachment]