Subject: Occasionally iptables blocks simply stop working
Posted by wsap on Wed, 08 Jan 2020 15:18:04 GMT
View Forum Message <> Reply to Message

Heya folks,

We've got one container out of hundreds where, every once in a while, it just stops actually blocking things with iptables, and we're having a hard time figuring out why.

We have at least a dozen containers running essentially the same software stack with CentOS 7 x64 and CSF (Firewall) with LFD (among other things running) and none of the others have encountered this issue.

Here's the timeline each time this occurs:

1. Notice higher than normal load on the container (2-3 rather than 0.5)
2. Check processes, see php-fpm processes using CPU. Monitor logs of that website and see xml-rpc attacks or wp-login attacks on WordPress sites. LFD has rules in place to detect both of these types of attacks on WordPress and add them to CSF.
3. Query CSF with csf -g {ip} -> CSF returns that the IP is either blocked in iptables directly or in an IPSET chain. (Note: if we disable IPSET in CSF's config and use only iptables, the result is the same, so I don't think this is specific to IPSET). Just to be sure, I also queried iptables directly using iptables -L -n | grep {ip} to confirm the IP is indeed listed there (when ipset is disabled) and the IP is definitely there and configured to DROP all packets from it.
4. Yet the bruteforce attack continues, despite clearly seeing LFD blocking the IP and finding the IP in iptables or IPSET's block list.

I've confirmed that the chain in iptables has a default policy of DROP. Running csf -r to reload the config does not resolve it.

Restarting the container always resolves it for some unknown period of time. Typically at least 24 hours - 1 week later the issue returns.

Recently I've had good luck resolving it without a reboot by running: systemctl restart network

The *one* major difference this container has over our others that are similarly configured is that it has a larger number of non-contiguous IPv4 addresses assigned.

While I don't know that this is specifically an OpenVZ problem, given that the entire network stack of any given container is emulated / provided by the host node in OpenVZ software code, it seems plausible.

Part of the problem is that we don't even know *when* it begins each time. Could be hours before, could be days before we detect the issue. Does anyone know where is best to look to find the source of the issue?

Thanks in advance for any guidance.