

Smart Proxy - Bug #31829

The proxies cached list of DHCP reservations gets out of sync with the DHCP server

02/05/2021 07:27 PM - Keith C

Status:	Need more information	
Priority:	Normal	
Assignee:		
Category:	DHCP	
Target version:		
Difficulty:		Fixed in Releases:
Triaged:	No	Found in Releases:
Bugzilla link:		Red Hat JIRA:
Pull request:		
Description		
<p>We have a foreman deployment that currently manages 20k physical hosts spread across a few dozen data centers (smart proxies). We often run into DHCP conflict errors when rebuilding a host's configuration (e.g. Building a host). The master asks the proxy if there is already a DHCP entry for this host, the proxy responds that their isn't, so the master tries to create one. The proxy then sets up the OMAPI calls to create the reservation, but DHCPd rejects the request because the DHCP record does actually exist. If I call the proxy's DHCP endpoint and ask it for that reservation, it returns None. If I then restart the foreman-proxy service and again ask the proxy to get the same reservation, it's returned as expected.</p> <p>So it seems that the proxy's cached view of the DHCP leases/reservations gets out of sync with reality until we reload the service.</p> <p>We're using stock ISC DHCP with proxy version 1.21.3.</p> <p>It's also worth noting that we're using the same DHCP pool for both static and dynamic leases. This was a mistake that we've fixed in the next generation of our deployment, but for the most part sharing the pool between static and dynamic has worked okay. After looking through the proxy code, I suspect this problem is being caused by the fact that the dynamic and static ranges are the same. Our provisioning flow goes like this:</p> <ol style="list-style-type: none">1. New host PXEs and gets a dynamic IP and boots into Discovery2. We convert the Discovered host into Provisioned, reserving the dynamic IP the host got when it first booted. <p>So because of this, there is a time period when both static and dynamic leases exist for the same host. I suspect that when the dynamic lease does finally expire, the proxy's DHCP observer treats this as the reservation being removed, so it removes the reservation from its internal cache. This is just a guess and could be a total red herring, but I haven't been able to find reports of anyone else running into this conflict issue.</p>		

History

#1 - 02/08/2021 09:53 AM - Lukas Zapletal

Hello, please use our forum for support requests, I don't see a particular bug isolated yet.

Proxy DHCP module uses inotify to detect when a lease file was modified, then it loads it into memory. If you use e.g. NFS inotify will not work at all. If you have 20k records it might take a while to load and parse the whole file. This is growing file it can grow quite a bit until it's squashed. ISC always squash the file during restart. I suggest you to measure how long it takes to parse that file, it's the `subnets_hosts_and_leases` method. I think it would be wise to add some timing into the debug log for everybody.

Yes, proxy should correctly treat expired leases, our parser sees that and marks those as released.

In regard to discovery lease, Foreman has a special code that should treat lease as non-conflict because we assume it could have been a discovered host.

I don't know what exactly do you mean by static reservations, Foreman only performs OMAPI calls and all these entries are dynamic only. Static reservations are done via configuration edit and Foreman cannot update those, it does not simply have permissions to do that.

#2 - 02/08/2021 04:28 PM - Keith C

Thanks for the response! By 'static reservation' I actually meant 'dynamic reservation' (as opposed to just a basic lease).

I haven't been able to trigger this problem in a test lab, it only shows up in our production deployments long after all the hosts have been provisioned. I only mentioned the size of our deployment to try to capture the fact that the problem is somewhat rare, but shows up regularly once the number of

managed hosts gets large enough. So the most succinct way I can describe this bug is:

1. Setup a bunch of Hosts in Foreman, including dynamic DHCP reservations for their primary and IPMI interfaces via foreman
2. Wait a few weeks...
3. Trigger a 'Rebuild Host Configuration' action on hosts until one of them fails with a DHCP conflict (conflicting reservation already exists)
4. Query the Smart Proxy API directly for the MAC or IP of the conflicting interface, observe that the smart proxy says there are no such reservations.
5. Restart the Foreman Proxy service
6. Again query the Smart Proxy API for the MAC or IP of the conflicting interface, observe that after the restart the proxy returns the reservation as expected.

So the reason I think this is a bug is that ONLY restarting the proxy service changes the behavior of the smart proxy API. However, I totally understand that this bug description isn't very actionable.

When I run into this again, I'll try touching the leases db file to force inotify to run and see if it helps. I'll also try timing the file parsing like you suggested.

Thanks again for the input. Any other suggestions for debugging further would be greatly appreciated!

#3 - 02/09/2021 07:38 AM - Lukas Zapletal

- Status changed from New to Need more information

Thanks, as soon as I get some more details I will do my best to help you. We don't have a 20k test environment, however ISC DHCP integration is clunky (inotify/parsing/omapi) and I am not surprised to see its limits. You can consider using dnsmasq (unofficial) plugins which is integrated in a different way (inotify used in the daemon to trigger actions, smart proxy just managing files in a directory).

Well let me know if you know something, I would be very interested in parsing speed of that huge file.