

Foreman - Feature #8824

Update dhcp record on host build state changed and disable PXE for hosts successfully built

01/02/2015 09:50 AM - Baptiste Agasse

Status:	Rejected	
Priority:	Normal	
Assignee:		
Category:	DHCP	
Target version:		
Difficulty:		Fixed in Releases:
Triaged:		Found in Releases:
Bugzilla link:		Red Hat JIRA:
Pull request:		

Description

The current behaviour in foreman for dhcp record updates is:

- When we create new host or update it network configuration, foreman tell to the dhcp smart proxy to update the 'host' record.

The created/updated record look like this:

```
host foobar.baz.example.com {
  dynamic;
  hardware ethernet 52:54:00:xx:xx:xx;
  fixed-address 192.168.0.90;
  supersede server.filename = "pxelinux.0";
  supersede server.next-server = c0:a8:00:03;
  supersede host-name = "foobar.baz.example.com";
}
```

So when we boot the 'foobar.baz.example.com', if we are in 'build mode' (with kickstart installation), the TFTP server have a specific version of the pxe config for this host with a link to the kickstart etc in order to install the host. When the install finish, foreman is notified that the host is successfully built and the pxe config on the TFTP server is modified to use a default pxe menu with only local boot. The dhcp record is not modified when the host goes from 'build mode' to 'normal mode' (or from 'normal mode' to 'build mode' again).

With this modification, the dhcp record is updated when the 'build' boolean of the host is modified too. If the host is in 'build mode' the default dhcp record is used. If the host is not in 'build mode', the dhcp record will be the following:

```
host foobar.baz.example.com {
  dynamic;
  hardware ethernet 52:54:00:xx:xx:xx;
  fixed-address 192.168.0.90;
  supersede server.filename = "";
  supersede server.next-server = "";
  supersede host-name = "foobar.baz.example.com";
}
```

This disable the PXE boot for the current host and then bypass the following steps:

- The dhcp server provide params
- The host download wmlinuz and initrd on the TFTP server
- The host print the PXE boot menu

The result is a faster boot when the host is installed and don't ask and download vmlinuz and initrd files on the network for local boot.

In more general usecase, if we use the discovery plugin or other default pxe configuration:

- If you boot the host in build mode, you get the 'install' pxe boot with the kicstart etc..

- If you boot the host in 'normal' mode (host successfully installed), pxe boot is disabled (the dhcp server don't provide you 'filename' and 'nextserver', so the boot process go ahead (commonly on your local disk)
- If you boot an unknown host, this host don't have a dhcp record, so it will use the defaults defined for it subnet or the global defaults (boot on discovery image if discovery plugin is configured)

Tested with foreman 1.7.1 and ISC DHCP server.

History

#1 - 01/02/2015 09:51 AM - Baptiste Agasse

You can find the code here: https://github.com/bagasse/foreman/tree/update_dhcp_record_on_host_build_state_changed
If you are interested, i can submit a pull request.

#2 - 01/04/2015 05:24 PM - Anonymous

While not against the change as such, you can edit the "PXELinux default local boot" template not to show the menu and without timeout boot from HDD (the PXE/TFTP process will still be there, but no menu and no 20 sec. delay). Would this be an option for you?

#3 - 01/05/2015 03:35 AM - Baptiste Agasse

Yes, i've discussed about this option with ohadlevy on irc. But for me modify the default pxe template is a workaround, it's not only about timeout. For me the build mode is a particular case, that need PXE, and i don't want that regular case (boot on local disk) be impacted by requirements of this particular case. I think this workflow (PXE install and then boot on local disk) is the main for foreman users that use PXE, but maybe i'm wrong.

For example, the vmlinuz and initrd files for CentOS 7 are about 40MB. I don't want that my hosts download 40MB each boot time, just to have a PXE boot that make the host boot on local disk.

Maybe we can enable this new workflow with a global setting in foreman, false by default for compatibility ?

Any comments are welcome.

#4 - 01/05/2015 03:43 AM - Lukas Zapletal

Hello, thanks for the contribution.

I don't really understand, what are you trying to avoid. Servers can PXEboot without loading Linux kernel into memory, this code is burnt into the NIC and therefore the PXELinux menu and stuff loads almost instantly. The only problem is the timeout which defaults to 20 seconds I believe. Can you describe your configuration and the reason why you need to load CentOS kernel and RAM disk into the memory? If you are loading it up, you must be PXEbooting it, therefore you can PXEboot directly the menu with the "Boot from hard drive" option instead of the kernel/image.

If this turns out to be valid case, please go ahead and file a Pull Request for that. The solution looks sane, but I'd go ahead with creating a global setting named like "Disable boot menu for provisioned hosts" that would be turned off by default. I think we should give existing users an option to opt-in (and in future if this turns out to be better option to opt-out).

#5 - 01/05/2015 05:18 AM - Baptiste Agasse

Ok when i checked the pxe config file generated for my test host, i made a typo in the mac address and i showed a pxe menu of another host that was in build mode... my assumption was therefore false. So modifying the timeout in the 'local boot default template will be ok. Sorry for the noise.

#6 - 01/05/2015 07:13 AM - Lukas Zapletal

- Status changed from New to Rejected

Okay your patch is still interesting but I am closing this ticket for now. If folks would like to extend this idea, please reopen and carry on.

Our default value is rather longer for those who want to see the menu or want to add more options there. If you like, please bring a discussion to the list about the default value. I also feel like it's maybe too long (5 seconds looks better to me). Hard to say.

Thanks.