

Installer - Bug #19081

node.rb failing with servers with large amount of facts (ips?)

03/29/2017 06:00 PM - David Ibarra

Status: New	
Priority: Normal	
Assignee:	
Category: Foreman modules	
Target version:	
Difficulty:	Fixed in Releases:
Triaged:	Found in Releases: 1.14.2
Bugzilla link:	Red Hat JIRA:
Pull request:	
Description	
Hey,	
I recently had some boxes come under my management that have a super large amount of IPs on them, like, 500-600 individual ip addresses bound to them.	
Foreman seems to ingest these facts really really slowly when /etc/puppet/node.rb is run. I end up getting a timeout after waiting around for 60 seconds:	
<pre>[root@foreman ~]# /etc/puppet/node.rb host1.hostname.com Could not send facts to Foreman: Net::ReadTimeout</pre>	
My solution was to raise this like so:	
<pre>--- /etc/puppet/node.rb 2017-03-29 16:53:12.755852699 -0500 +++ /etc/puppet/node.rb.new 2017-03-29 16:47:06.660634862 -0500 @@ -107,6 +107,7 @@ def initialize_http(uri) res = Net::HTTP.new(uri.host, uri.port) res.use_ssl = uri.scheme == 'https' + res.read_timeout = 180 if res.use_ssl? if SETTINGS[:ssl_ca] && !SETTINGS[:ssl_ca].empty? res.ca_file = SETTINGS[:ssl_ca]</pre>	
I know this is inelegant, first of all it should probably use SETTINGS[:enc_read_timeout] or something easier to configure, but closer to the issue, I'm not sure why it takes so long for Foreman to ingest these facts. I tried tracing through the code but wasn't really successful in figuring it out.	
Related issues:	
Related to Foreman - Bug #20929: Huge table "fact_names" makes Foreman crawl	Closed 09/14/2017

History

#1 - 03/30/2017 02:56 AM - Dominic Cleal

- Project changed from Foreman to Installer

- Category changed from Facts to Foreman modules

The script already has a :timeout setting, which is set to read_timeout (and open_timeout in the latest version) when performing fact uploads. Shouldn't that change the same thing?

#2 - 03/30/2017 11:23 AM - David Ibarra

Dominic Cleal wrote:

The script already has a :timeout setting, which is set to read_timeout (and open_timeout in the latest version) when performing fact uploads. Shouldn't that change the same thing?

Sort of. I don't claim to have advanced knowledge of the foreman code base, or even Ruby, but it seems that `:timeout` for `node.rb` controls a timeout for the block of code that gets the classification of the server, not the `read_timeout` for the `Net::HTTP` client that pushes the facts into the foreman API.

#3 - 04/14/2017 05:41 PM - David Ibarra

I think I found the cause of the slowness, it seems like it's ActiveRecord:

`app/models/host/base.rb`:

```
if iface.new_record? || iface.changed?
  logger.debug "Saving #{name} NIC for host #{self.name}"
  result = iface.save

  unless result
    logger.warn "Saving #{name} NIC for host #{self.name} failed, skipping because:"
    iface.errors.full_messages.each { |e| logger.warn " #{e}" }
  end

  result
end
```

If I disable activerecord callbacks (`iface.save` to `iface.save(:validate => false)`) this thing goes way way faster. I'm talking like, we're going from 5 minutes to import these facts, to 14 seconds.

#4 - 12/19/2017 01:02 PM - Ewoud Kohl van Wijngaarden

- Related to Bug #20929: Huge table "fact_names" makes Foreman crawl added