Installer - Feature #18821

Set PassengerMaxRequests and PassengerMemoryLimit to reasonable values

03/07/2017 09:02 AM - Lukas Zapletal

Status:	Closed		
Priority:	Normal		
Assignee:	Lukas Zapletal		
Category:	Foreman modules		
Target version:	1.15.0		
Difficulty:		**Fixed in Releases:**	
Triaged:		**Found in Releases:**	
Bugzilla link:		**Red Hat JIRA:**	
Pull request:	https://github.com/theforeman/foreman-installer/pull/221		

Description

We leave these values unconfigured which means no restarting is performed. Ruby VM is known to grow on memory and due nature of it's memory model, it never returns memory back to OS and due fragmentation it grows slowly. The only way out on CRuby (MRI) is to restart processes after some time, it should not happend more than once per day as a generic rule of thumb.

Now, Passenger provides tools to do this automatically, it can restart after N served requests and/or after memory goes above threshold. We should be able to setup reasonable default values (100,000 requests and 1 GB memory per instance) to prevent OOM kills which happens on production installs.

https://www.phusionpassenger.com/library/config/apache/reference/#passengermemorylimit

https://www.phusionpassenger.com/library/config/apache/reference/#passengermaxrequests

Associated revisions

Revision b570b53b - 03/09/2017 06:18 AM - Lukas Zapletal

Fixes #18821 - set PassengerMaxRequests to keep processes healthy

History

#1 - 03/08/2017 04:14 AM - Ewoud Kohl van Wijngaarden

Given the PassengerMemoryLimit has a big warning that it's enterprise only we can't use that. For the other setting they mention 'This option should be considered as a workaround for misbehaving applications. It is advised that you fix the problem in your application rather than relying on this option as a measure to avoid memory leaks.' so setting this by default doesn't feel right.

As a workaround you can add the following to /etc/foreman-installer/custom-hiera.yaml

```
---
apache::mod::passenger::passenger_max_requests: 100000
```

This will set it globally for the server, but currently there's no parameter to easily set for the single vhost.

#2 - 03/08/2017 12:38 PM - Ewoud Kohl van Wijngaarden

I just talked to Timo and they deploy it with max requests set to 10000, at least on their puppet masters. Please submit a patch to https://github.com/theforeman/foreman-installer/blob/develop/config/foreman.hiera/tuning.yaml to fix it for Foreman. I don't know if Katello will also use it then but that should be easy to verify.

#3 - 03/09/2017 03:48 AM - Lukas Zapletal

Ewoud, this is not a memory leak, every single Ruby application eventually needs restarting, if it allocates some memory. There was a great talk at FOSDEM about this. Ruby will effectively never return allocated memory to OS, so GC will not help at all and the only way to free memory is to terminate the process.

I don't understand why you accept 100000 requests limit but not like 1 GB top limit, which is IMHO much more useful than 100000 which is little bit artificial (while we know that Foreman w/ Katello takes about 400 MB so 1 GB is sane default to initiate restart).

#4 - 03/09/2017 04:01 AM - Ewoud Kohl van Wijngaarden

The reason I wouldn't accept the memory limit is that as I read it, it's not available in the open source release, only in the enterprise release. That means it's not available to us.

#5 - 03/09/2017 06:17 AM - Lukas Zapletal

Oh really is it? Is this cripple-ware? Jeeez. I had no idea, I live in a perfect world I guess :-)

I will make the patch then.

#6 - 03/09/2017 06:19 AM - The Foreman Bot

- Status changed from New to Ready For Testing

- Assignee set to Lukas Zapletal

- Pull request https://github.com/theforeman/foreman-installer/pull/221 added

#7 - 03/09/2017 01:01 PM - Lukas Zapletal

- Status changed from Ready For Testing to Closed

- % Done changed from 0 to 100

Applied in changeset [b570b53b8e7df5cb0fdff30029a7e0c1a4c60dff](#).

#8 - 04/03/2017 03:02 AM - Dominic Cleal

- translation missing: en.field_release set to 209