

## Foreman - Bug #11834

### host status migration is very slow

09/15/2015 01:36 PM - Ohad Levy

<b>Status:</b> Closed	
<b>Priority:</b> Normal	
<b>Assignee:</b> Marek Hulán	
<b>Category:</b> DB migrations	
<b>Target version:</b> 1.10.1	
<b>Difficulty:</b>	<b>Fixed in Releases:</b>
<b>Triaged:</b>	<b>Found in Releases:</b> Nightly
<b>Bugzilla link:</b>	<b>Red Hat JIRA:</b>
<b>Pull request:</b> <a href="https://github.com/foreman/foreman/pull/2860">https://github.com/foreman/foreman/pull/2860</a>	
<b>Description</b>	
On a largish db, calculating the host status takes a whole lot of time.  it seems that many queries are repeating them selfs are simply could be eager loaded for the migration.	
<b>Related issues:</b>	
Related to Foreman - Feature #10782: Add global status for hosts	<b>Closed</b> <b>06/11/2015</b>

#### Associated revisions

##### Revision 553665c2 - 11/23/2015 03:10 AM - Marek Hulán

Fixes #11834 - improve host status migration performance

##### Revision 7922ad2f - 01/19/2016 08:10 AM - Marek Hulán

Fixes #11834 - improve host status migration performance

(cherry picked from commit 553665c2a458b949d85c45aac9fea33271a2a9a0)

#### History

##### #1 - 09/15/2015 01:37 PM - Ohad Levy

- Related to Feature #10782: Add global status for hosts added

##### #2 - 09/16/2015 02:54 AM - Dominic Cleal

- Assignee deleted (Marek Hulán)

- translation missing: en.field\_release set to 63

Can you provide some logs to show which queries?

##### #3 - 09/16/2015 03:58 AM - Ohad Levy

1. we do host.each, not sure if we can avoid it, but we could try to loop in batches and load associated data.
2. the status calculation is done per host, which might not be avoidable, but I would like to see if we can wrap it in a host collection instead of a single host to avoid many n+1 queries.

##### #4 - 10/23/2015 06:43 AM - The Foreman Bot

- Status changed from New to Ready For Testing

- Pull request <https://github.com/foreman/foreman/pull/2860> added

- Pull request deleted ()

##### #5 - 10/23/2015 06:43 AM - Marek Hulán

I tried to fix some (n+1) queries and do the processing in batches (batch size didn't seem to have much effect), I don't see any major improvement. On the other hand, my DB is quite small.

**#6 - 11/23/2015 03:15 AM - Dominic Cleal**

- Assignee set to Marek Hulán
- translation missing: en.field\_release changed from 63 to 104

**#7 - 11/23/2015 04:01 AM - Marek Hulán**

- Status changed from Ready For Testing to Closed
- % Done changed from 0 to 100

Applied in changeset [553665c2a458b949d85c45aac9fea33271a2a9a0](#).