

RPM Support - Issue #630

RPM unit types have redundant search indices specified

02/28/2015 11:45 PM - rbarlow

Status: CLOSED - DUPLICATE	Start date:
Priority: Normal	Due date:
Assignee:	Estimated time: 0:00 hour
Category:	
Sprint/Milestone:	
Severity: 1. Low	Groomed: No
Version: 2.4.0	Sprint Candidate: No
Platform Release:	Tags: Pulp 2
OS:	Sprint:
Triaged: Yes	Quarter:

Description

Description of problem:
The RPM unit type JSON file specifies the first field of the unit key as a search index on multiple types[0]. For example, RPM lists "name" as a search index and also as the first element of the unit key. This is redundant since the unit key causes a uniqueness constraint on the collection that effectively gives us a free "name" index. The harm in doing this is that MongoDB will use more RAM and will make writes slower since it has to update two indices needlessly.

Version-Release number of selected component (if applicable):
2.4.3-1

How reproducible:
Every time.

Steps to Reproduce:
1. Look at /usr/lib/pulp/plugins/types/rpm_support.json

or

1. Use Mongo's shell to inspect the installed indices on the various unit collections.

Actual results:
There are redundant search indices.

Expected results:
There should not be any redundant search indices.

Additional info:
[0] https://github.com/pulp/pulp_rpm/blob/2.4-release/plugins/types/rpm_support.json

+ This bug was cloned from [Bugzilla Bug #1169503](#) +

Related issues:
Is duplicate of RPM Support - Issue #1476: The RPM models are indexing a lot ... **CLOSED - WONTFIX**

History

#1 - 02/28/2015 11:45 PM - rbarlow

It may also be worth reconsidering all the fields we are indexing for searching while we fix this bug. For example, some of the units are indexing "arch", which will likely not be useful and will waste resources. The decision on what to index and what not to index can be tricky, but try to think about the balance of what would be useful to end users vs. how much RAM/write time each index costs.

+ This comment was cloned from [Bugzilla #1169503 comment 1](#) +

#2 - 03/20/2015 08:13 PM - bmbouter

- Severity changed from Low to 1. Low

#3 - 02/23/2016 10:21 PM - rbarlow

- Is duplicate of Issue #1476: *The RPM models are indexing a lot of fields that don't make sense to have indices added*

#4 - 02/23/2016 10:23 PM - rbarlow

- Status changed from NEW to CLOSED - DUPLICATE

The other bug report offers better details.

#5 - 04/15/2019 11:07 PM - bmbouter

- Tags Pulp 2 added