

RPM Support - Issue #1565

Pulp does not publish more than once

01/21/2016 10:16 PM - dkliban@redhat.com

Status:	CLOSED - CURRENTRELEASE	Start date:	
Priority:	High	Due date:	
Assignee:	dkliban@redhat.com	Estimated time:	0:00 hour
Category:		Groomed:	No
Sprint/Milestone:		Sprint Candidate:	No
Severity:	3. High	Tags:	Pulp 2
Version:		Sprint:	
Platform Release:	2.8.0	Quarter:	
OS:			
Triaged:	Yes		

Description

The simplest way to see this is to create a repository, publish it, and then sync it (with auto publish).

repo.json

```
{"id": "test-repo", "display_name": "test-repo", "importer_type_id": "yum_importer", "importer_config": {"feed": "https://repos.fedorapeople.org/pulp/pulp/demo_repos/zoo/", "ssl_ca_cert": null, "ssl_client_cert": null, "ssl_client_key": null}, "notes": {"_repo-type": "rpm-repo"}, "distributors": [{"distributor_type_id": "yum_distributor", "distributor_config": {"relative_url": "test-repo", "http": true, "https": true, "protected": true}, "auto_publish": true, "distributor_id": "test-repo"}, {"distributor_type_id": "export_distributor", "distributor_config": {"http": false, "https": false}, "auto_publish": false, "distributor_id": "export_distributor"}]}
```

publish.json

```
{"id": "test-repo", "override_config": null}
```

sync.json

```
{"override_config": {"num_threads": 4, "validate": true}}
```

Steps to reproduce:

```
$ http --verify no --auth admin:admin POST https://localhost/pulp/api/v2/repositories/ < repo.json
```

```
$ http --verify no --auth admin:admin POST https://localhost/pulp/api/v2/repositories/test-repo/actions/publish/ < publish.json
```

```
$ http --verify no --auth admin:admin POST https://localhost/pulp/api/v2/repositories/test-repo/actions/sync/ < sync.json
```

Associated revisions

Revision 9f0c6ecb - 01/25/2016 09:30 PM - dkliban@redhat.com

Adds validation to the UTCDateTimeField for mongoengine

Due to a bug in datetime utils, we can't rely on mongonegine to convert a string representation of a timestamp with timezone represented as Z on end. As a result we now require the date to be supplied as a datetime object with tzinfo set.

<https://pulp.plan.io/issues/1565> closes #1565

Revision 9f0c6ecb - 01/25/2016 09:30 PM - dkliban@redhat.com

Adds validation to the UTCDateTimeField for mongoengine

Due to a bug in datetime utils, we can't rely on mongonegine to convert a string representation of a timestamp with timezone represented as Z on end. As a result we now require the date to be supplied as a datetime object with tzinfo set.

History

#1 - 01/21/2016 10:36 PM - dkliban@redhat.com

- Description updated

#2 - 01/22/2016 04:28 PM - dkliban@redhat.com

- Description updated

#3 - 01/22/2016 05:09 PM - jortel@redhat.com

- Assignee set to dkliban@redhat.com

- Priority changed from Normal to High

- Triaged changed from No to Yes

#4 - 01/22/2016 10:45 PM - dkliban@redhat.com

Pulp seems to expect that the date on the machine is in UTC timezone. I was able to resolve this problem by changing the timezone from EST to UTC.

```
sudo rm /etc/localtime
sudo ln -s /usr/share/zoneinfo/UTC /etc/localtime
```

I am still determining which date in Pulp codebase is wrong.

#5 - 01/24/2016 11:13 PM - dkliban@redhat.com

The problem stems from a bug[0] in the `dateutil.parser.parse()`. This method is used by mongoengine to parse a string to create a python datetime object. It seems that `parse()` does not recognize the timezone part of the string and makes the assumption that it is in localtime. This works fine when a machine has its local time set to UTC. When timezone on a machine is anything other than UTC, handing a UTC date to `parse()` causes the date to be converted to UTC from local time.

In this case[1], when we are saving the last published date, we first generate it as a string that looks like "2016-01-24T20:46:25Z". We then assign the value to 'last_publish' field of the Distributor object. During the `save()` of the distributor, the `parse()` method is used to convert the string to a datetime object. The Z part of the timezone is interpreted as being local time instead of UTC. So the database ends up storing a date that's 5 hours in the future (in the EST case).

[0] <https://github.com/dateutil/dateutil/issues/70>

[1] <https://github.com/pulp/pulp/blob/master/server/pulp/server/controllers/repository.py#L1000-L1002>

#6 - 01/25/2016 09:31 PM - dkliban@redhat.com

- Status changed from NEW to POST

<https://github.com/pulp/pulp/pull/2361>

#7 - 01/25/2016 09:32 PM - dkliban@redhat.com

- Status changed from POST to MODIFIED

- % Done changed from 0 to 100

Applied in changeset [pulp:pulp9f0c6ecb7d943abc8fefe4c3c131796c90e9cdee](https://pulp.plan.io/changeset/pulp:pulp9f0c6ecb7d943abc8fefe4c3c131796c90e9cdee).

#8 - 02/11/2016 10:26 PM - rbarlow

- Status changed from MODIFIED to 5

#9 - 03/14/2016 10:25 PM - pthomas@redhat.com

- Status changed from 5 to 6

verified

```
root@mgmt5 ~]#
```

```
[root@mgmt5 ~]# http --verify no --auth admin:admin POST https://localhost/pulp/api/v2/repositories/ < repo.json
/usr/lib/python2.7/site-packages/urllib3/connectionpool.py:769: InsecureRequestWarning: Unverified HTTPS request is being made. Adding certificate
verification is strongly advised. See: https://urllib3.readthedocs.org/en/latest/security.html
InsecureRequestWarning)
HTTP/1.1 201 CREATED
Connection: Keep-Alive
```

Content-Length: 313
Content-Type: application/json; charset=utf-8
Date: Mon, 14 Mar 2016 21:24:47 GMT
Keep-Alive: timeout=5, max=10000
Location: <https://localhost/pulp/api/v2/repositories/test-repo/>
Server: Apache/2.4.6 (Red Hat Enterprise Linux) OpenSSL/1.0.1e-fips mod_wsgi/3.4 Python/2.7.5

```
{
  "_href": "/pulp/api/v2/repositories/test-repo/",
  "_id": {
    "$oid": "56e72c1fbc1dbc60953fd076"
  },
  "_ns": "repos",
  "content_unit_counts": {},
  "description": null,
  "display_name": "test-repo",
  "id": "test-repo",
  "last_unit_added": null,
  "last_unit_removed": null,
  "notes": {
    "_repo-type": "rpm-repo"
  },
  "scratchpad": {}
}
```

```
[root@mgmt5 ~]# http --verify no --auth admin:admin POST https://localhost/pulp/api/v2/repositories/test-repo/actions/publish/ < publish.json
/usr/lib/python2.7/site-packages/urllib3/connectionpool.py:769: InsecureRequestWarning: Unverified HTTPS request is being made. Adding certificate
verification is strongly advised. See: https://urllib3.readthedocs.org/en/latest/security.html
InsecureRequestWarning)
HTTP/1.1 202 ACCEPTED
Connection: Keep-Alive
Content-Length: 172
Content-Type: application/json; charset=utf-8
Date: Mon, 14 Mar 2016 21:24:52 GMT
Keep-Alive: timeout=5, max=10000
Server: Apache/2.4.6 (Red Hat Enterprise Linux) OpenSSL/1.0.1e-fips mod_wsgi/3.4 Python/2.7.5
```

```
{
  "error": null,
  "result": null,
  "spawned_tasks": [
    {
      "_href": "/pulp/api/v2/tasks/54f41813-b4bb-41c7-8133-9ac2e576935c/",
      "task_id": "54f41813-b4bb-41c7-8133-9ac2e576935c"
    }
  ]
}
```

```
[root@mgmt5 ~]# http --verify no --auth admin:admin POST https://localhost/pulp/api/v2/repositories/test-repo/actions/sync/ < sync.json
/usr/lib/python2.7/site-packages/urllib3/connectionpool.py:769: InsecureRequestWarning: Unverified HTTPS request is being made. Adding certificate
verification is strongly advised. See: https://urllib3.readthedocs.org/en/latest/security.html
InsecureRequestWarning)
HTTP/1.1 202 ACCEPTED
Connection: Keep-Alive
Content-Length: 172
Content-Type: application/json; charset=utf-8
Date: Mon, 14 Mar 2016 21:25:04 GMT
Keep-Alive: timeout=5, max=10000
Server: Apache/2.4.6 (Red Hat Enterprise Linux) OpenSSL/1.0.1e-fips mod_wsgi/3.4 Python/2.7.5
```

```
{
  "error": null,
  "result": null,
  "spawned_tasks": [
    {
      "_href": "/pulp/api/v2/tasks/0aa72df6-5719-436b-a0da-bdfb4bda2d53/",
      "task_id": "0aa72df6-5719-436b-a0da-bdfb4bda2d53"
    }
  ]
}
```

#10 - 03/23/2016 07:48 PM - dkliban@redhat.com

- Status changed from 6 to CLOSED - CURRENTRELEASE

#12 - 04/15/2019 10:37 PM - bmbouter

- Tags Pulp 2 added