

Pulp - Story #9304

As a user, I can configure https proxy using https://

08/27/2021 02:00 PM - ttereshc

Status:	CLOSED - NOTABUG	Start date:	
Priority:	High	Due date:	
Assignee:	bmbouter	% Done:	0%
Category:		Estimated time:	0:00 hour
Sprint/Milestone:		Tags:	Katello
Platform Release:		Sprint:	Sprint 106
Groomed:	No	Quarter:	
Sprint Candidate:	No		

Description

If one configures https proxy, sync task fails with the following error.

```
pulp_tasks:
```

```
- pulp_href: "/pulp/api/v3/tasks/c3aadbfef-2b3d-449d-8e65-9666951476d2/"
  pulp_created: '2021-08-16T11:20:14.160+00:00'
  state: failed
  name: pulp_rpm.app.tasks.synchronizing.synchronize
  logging_cid: 0f189218-3ac5-4a3f-b183-0d6b916aa4f2
  started_at: '2021-08-16T11:20:14.249+00:00'
  finished_at: '2021-08-16T11:20:14.374+00:00'
  error:
    traceback: |2
      File "/usr/lib/python3.6/site-packages/pulpcore/tasking/pulpcore_worker.py", line 272, in
      _perform_task
        result = func(*args, **kwargs)
      File "/usr/lib/python3.6/site-packages/pulp_rpm/app/tasks/synchronizing.py", line 363, in
      synchronize
        remote_url = fetch_remote_url(remote)
      File "/usr/lib/python3.6/site-packages/pulp_rpm/app/tasks/synchronizing.py", line 242, in
      fetch_remote_url
        get_repomd_file(remote, normalized_remote_url)
      File "/usr/lib/python3.6/site-packages/pulp_rpm/app/tasks/synchronizing.py", line 200, in
      get_repomd_file
        return downloader.fetch()
      File "/usr/lib/python3.6/site-packages/pulpcore/download/base.py", line 176, in fetch
        return done.pop().result()
      File "/usr/lib/python3.6/site-packages/pulpcore/download/http.py", line 258, in run
        return await download_wrapper()
      File "/usr/lib/python3.6/site-packages/backoff/_async.py", line 133, in retry
        ret = await target(*args, **kwargs)
      File "/usr/lib/python3.6/site-packages/pulpcore/download/http.py", line 256, in download_w
      rapper
        return await self._run(extra_data=extra_data)
      File "/usr/lib/python3.6/site-packages/pulp_rpm/app/downloaders.py", line 92, in _run
        url, proxy=self.proxy, proxy_auth=self.proxy_auth, auth=self.auth
      File "/usr/lib64/python3.6/site-packages/aiohttp/client.py", line 1117, in __aenter__
        self._resp = await self._coro
      File "/usr/lib64/python3.6/site-packages/aiohttp/client.py", line 513, in _request
        traces=traces,
      File "/usr/lib64/python3.6/site-packages/aiohttp/client_reqrep.py", line 311, in __init__
        self.update_proxy(proxy, proxy_auth, proxy_headers)
      File "/usr/lib64/python3.6/site-packages/aiohttp/client_reqrep.py", line 551, in update_pr
      oxy
        raise ValueError("Only http proxies are supported")
    description: Only http proxies are supported
  worker: "/pulp/api/v3/workers/204569e0-9f35-48ca-9e22-5cbfd8828dcc/"
  child_tasks: []
```

```
progress_reports: []
created_resources: []
reserved_resources_record:
- "/pulp/api/v3/repositories/rpm/rpm/abc9dd0c-d198-433c-af59-47b40de82db6/"
- "/pulp/api/v3/remotes/rpm/rpm/d86a3ede-74f5-4f11-b979-34e9f0e6dd27/"
create_version: true
task_groups: []
poll_attempts:
  total: 1
  failed: 1
```

aiohttp doesn't support https proxies, specified with https://.

https://docs.aiohttp.org/en/stable/client_advanced.html#proxy-support

"aiohttp supports plain HTTP proxies and HTTP proxies that can be upgraded to HTTPS via the HTTP CONNECT method.

aiohttp does not support proxies that must be connected to via https://."

HTTP CONNECT method <https://www.ietf.org/rfc/rfc2817.txt>

Pulp 2 supported it, so the lack of this feature is seen as a regression.

Related issues:

Related to Pulp - Task #7921: Improve documentation on types of proxies suppo...

NEW

History

#1 - 08/27/2021 02:04 PM - ttereshc

It seems to be a low priority on the aiohttp list to complete <https://github.com/aio-libs/aiohttp/issues/845#issuecomment-313141691>. But reading the whole thread, it seems that they are not against a contribution, and the main problem is a test setup.

#2 - 08/30/2021 07:53 PM - bmbouter

- Status changed from NEW to ASSIGNED

- Assignee set to bmbouter

#3 - 09/01/2021 10:00 PM - bmbouter

I setup squid on a dev box and configured it to also support https_port like this:

1. Install with `sudo dnf install squid -y`
2. Make certificates with these commands (and use your hostname):

```
sudo openssl genrsa -out /etc/ssl/key.pem 4096
sudo openssl req -new -x509 -key /etc/ssl/key.pem -out /etc/ssl/cert.pem -days 1826
```

1. Add this line to the top of your `/etc/squid/squid.conf`: `https_port 3130 tls-cert=/etc/ssl/cert.pem tls-key=/etc/ssl/key.pem`

Then do some testing with curl.

Make a no proxy request: `curl -v http://example.com` Make a request going through the http-based proxy: `curl -v --proxy https://localhost:3128 http://example.com` Make a request going through the https proxy. `curl -v --proxy-insecure --proxy https://localhost:3130 http://example.com`

You have to use the `--proxy-insecure` because the certificate squid is using is self-signed. I couldn't easily get a request to `https://example.com` working, but that's not important for this bugfix because the proxy itself is TLS enabled.

Then have `pulp_file` sync through the http based proxy with:

```
#!/bin/bash
```

```
set -ev
```

```
repo_name="repo$RANDOM"
remote_name="remote$RANDOM"
distribution_name="distribution$RANDOM"
base_path="$RANDOM"
url="http://fixtures.pulpproject.org/file/PULP_MANIFEST"
proxy_url="http://pulp3-source-fedora33.fedora.example.com:3128/"
```

```
pulp file remote create --name $remote_name --url $url --proxy-url $proxy_url
pulp file repository create --name $repo_name --remote $remote_name
pulp file repository sync --name $repo_name
```

```
pulp file publication create --repository $repo_name
pulp file distribution create --base-path $base_path --repository $repo_name --name $distribution_name

echo "Created, synced, published repo $repo_name at distribution base path '$base_path'"
```

Then reproduce the issue by having it sync through the HTTPS proxy:

```
#!/bin/bash

set -ev

repo_name="repo$RANDOM"
remote_name="remote$RANDOM"
distribution_name="distribution$RANDOM"
base_path="$RANDOM"
url="http://fixtures.pulpproject.org/file/PULP_MANIFEST"
proxy_url="https://pulp3-source-fedora33.fedora.example.com:3130/"

pulp file remote create --name $remote_name --url $url --proxy-url $proxy_url
pulp file repository create --name $repo_name --remote $remote_name
pulp file repository sync --name $repo_name
pulp file publication create --repository $repo_name
pulp file distribution create --base-path $base_path --repository $repo_name --name $distribution_name

echo "Created, synced, published repo $repo_name at distribution base path '$base_path'"
```

It should show you: Error: Task /pulp/api/v3/tasks/ba5ecbf2-f7a8-48db-89e3-2c50af8fb8a6/ failed: 'Only http proxies are supported'

#4 - 09/10/2021 04:07 PM - bmbouter

Here's an update of the key points I've learned so far.

- I've been mostly discussing this on this aiohttp PR <https://github.com/aio-libs/aiohttp/pull/5992>. Lots of details there
- aiohttp doesn't do TLS in TLS correctly because Python doesn't do it correctly <https://bugs.python.org/issue37179>. This is close to being resolved and will be fixed in either 3.10 (maybe?) or 3.11 (likely).
- Until this bug is fixed aiohttp and other web frameworks would have to carry a workaround. Other ones like httpx do carry this.

So there are two options I see:

1. Continue to try to fix aiohttp's TLS handshake workaround code
2. Accept this isn't working until aiohttp fixes it properly after Python fixes asyncio

#5 - 09/10/2021 08:16 PM - rchan

- Sprint set to Sprint 105

#6 - 09/10/2021 09:52 PM - bmbouter

Here's another way to setup the proxy using a python-based proxy proxy.py

```
# Make your keys
sudo openssl genrsa -out /etc/ssl/key.pem 4096
sudo openssl req -new -x509 -key /etc/ssl/key.pem -out /etc/ssl/cert.pem -days 1826

# Install proxy and run it
pip install proxy
proxy --hostname 0.0.0.0 --cert-file /etc/ssl/cert.pem --key-file /etc/ssl/key.pem

# Test it with curl in another tty
curl -v --proxy-insecure --proxy https://127.0.0.1:8899 https://example.com
```

#7 - 09/23/2021 07:10 PM - dalley

- Related to Task #7921: Improve documentation on types of proxies supported added

#8 - 09/23/2021 11:54 PM - rchan

- Sprint changed from Sprint 105 to Sprint 106

#9 - 09/28/2021 08:09 PM - dalley

- Priority changed from Normal to High

#10 - 09/29/2021 10:25 PM - bmbouter

- Status changed from ASSIGNED to CLOSED - NOTABUG

After working with the aiohttp maintainers, this can only be fixed in aiohttp directly (PRs links below). There isn't a code change that Pulp can make to resolve it. As such, I'm closing as NOTABUG given that the bug isn't actually in Pulp.

<https://github.com/aio-lib/aiohttp/pull/5992/files><https://github.com/aio-lib/aiohttp/pull/6002/files>