

## Arvados - Bug #15944

### [Arvbox] Publicdev and keep issues

12/20/2019 08:35 PM - Lucas Di Pentima

<b>Status:</b>	Resolved	<b>Start date:</b>	01/24/2020
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assigned To:</b>	Lucas Di Pentima	<b>% Done:</b>	100%
<b>Category:</b>		<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	2020-01-29 Sprint		
<b>Description</b>			
When using arvbox in publicdev mode, access attempts to Keep fail with 'connection refused' messages. Maybe they're not being setup to listen to the public IP:			
<pre>2019-12-20T20:30:10.939000582Z crunch-run dev started 2019-12-20T20:30:10.939030767Z Executing container 'xzky2-dz642-ee96941wicte9e' 2019-12-20T20:30:10.939041998Z Executing on host '72f6da3e099a' 2019-12-20T20:30:11.158294161Z Fetching Docker image from collection 'dc979b1f82e04135f13f9f214c02e866+261' 2019-12-20T20:30:11.200995059Z Using Docker image id 'sha256:88a16dfa7bee916aea5d0d400a2a78da5dfc6aa9eeebf4f78e8ef33ee096864b' 2019-12-20T20:30:11.201982367Z Loading Docker image from keep 2019-12-20T20:30:11.203884681Z error in Run: While loading container image: While loading container image into Docker: error during connect: Post http://%2Fvar%2Frun%2Fdocker.sock/v1.21/images/load?quiet=1: GET 3cfb6d856d2b98e3d6a586edda419b8d+67108864+A14e8a1a993b9d71e0f88023079f2eeabcd941f49@5e0fa453 failed: [https://192.168.1.125:25101/3cfb6d856d2b98e3d6a586edda419b8d+67108864+A14e8a1a993b9d71e0f88023079f2eeabcd941f49@5e0fa453: Get https://192.168.1.125:25101/3cfb6d856d2b98e3d6a586edda419b8d+67108864+A14e8a1a993b9d71e0f88023079f2eeabcd941f49@5e0fa453: dial tcp 192.168.1.125:25101: connect: connection refused https://192.168.1.125:25101/3cfb6d856d2b98e3d6a586edda419b8d+67108864+A14e8a1a993b9d71e0f88023079f2eeabcd941f49@5e0fa453: Get https://192.168.1.125:25101/3cfb6d856d2b98e3d6a586edda419b8d+67108864+A14e8a1a993b9d71e0f88023079f2eeabcd941f49@5e0fa453: dial tcp 192.168.1.125:25101: connect: connection refused https://192.168.1.125:25101/3cfb6d856d2b98e3d6a586edda419b8d+67108864+A14e8a1a993b9d71e0f88023079f2eeabcd941f49@5e0fa453: Get https://192.168.1.125:25101/3cfb6d856d2b98e3d6a586edda419b8d+67108864+A14e8a1a993b9d71e0f88023079f2eeabcd941f49@5e0fa453: dial tcp 192.168.1.125:25101: connect: connection refused] 2019-12-20T20:30:11.225938552Z error in CaptureOutput: error scanning files to copy to output: open .arvados#collection: no such file or directory 2019-12-20T20:30:11.257592767Z Cancelled 2019-12-20T20:30:11.287732711Z error in CommitLogs: error saving log collection: error creating log manifest: Could not write sufficient replicas: [0] Put https://192.168.1.125:25101/7fe0ef0a324d0e77d6eb2398449aa05f: dial tcp 192.168.1.125:25101: 2019-12-20T20:30:11.386394909Z crunch-run finished</pre>			
<b>Subtasks:</b>			
Task # 15976: Review 15944-arvbox-publicdev-fix			<b>Resolved</b>

### Associated revisions

#### Revision c88bd214 - 01/28/2020 06:51 PM - Lucas Di Pentima

Merge branch '15944-arvbox-publicdev-fix'

Closes #15944

Arvados-DCO-1.1-Signed-off-by: Lucas Di Pentima <[lucas@di-pentima.com.ar](mailto:lucas@di-pentima.com.ar)>

### History

#### #1 - 12/26/2019 05:01 PM - Lucas Di Pentima

- Target version set to 2020-01-02 Sprint
- Assigned To set to Lucas Di Pentima
- Status changed from New to In Progress

#### #2 - 12/27/2019 03:48 PM - Lucas Di Pentima

Just for the record:

I've already found the missing published ports (keepproxy & arv-git-httpd) and now I'm getting a "loop detected" error when trying to upload something to Keep. It seems that for some reason keepproxy redirects the request to itself, but this doesn't happen when running arvbox on "dev" mode.

Below are the diffs of nginx.conf and cluster\_config.yml files between "dev" and "publicdev" modes:

## cluster\_config.yml

```
--- cluster_config.yml.dev      2019-12-26 20:46:52.409158517 -0300
+++ cluster_config.yml.publicdev 2019-12-26 20:37:57.816523439 -0300
@@ -7,13 +7,13 @@
     InternalURLs:
       "http://localhost:8004": {}
   Workbench1:
-   ExternalURL: "https://172.17.0.2:443"
+   ExternalURL: "https://192.168.1.125:443"
   Workbench2:
-   ExternalURL: "https://172.17.0.2:3001"
+   ExternalURL: "https://192.168.1.125:3001"
   SSO:
-   ExternalURL: "https://172.17.0.2:8900"
+   ExternalURL: "https://192.168.1.125:8900"
   Keepproxy:
-   ExternalURL: "https://172.17.0.2:25101"
+   ExternalURL: "https://192.168.1.125:25101"
     InternalURLs:
       "http://localhost:25100": {}
   Keepstore:
@@ -21,29 +21,29 @@
       "http://localhost:25107": {}
       "http://localhost:25108": {}
   Websocket:
-   ExternalURL: "wss://172.17.0.2:8002/websocket"
+   ExternalURL: "wss://192.168.1.125:8002/websocket"
     InternalURLs:
       "http://localhost:8005": {}
   GitSSH:
-   ExternalURL: "ssh://git@172.17.0.2:"
+   ExternalURL: "ssh://git@192.168.1.125:"
   GitHTTP:
     InternalURLs:
       "http://localhost:9001/": {}
-   ExternalURL: "https://172.17.0.2:9000/"
+   ExternalURL: "https://192.168.1.125:9000/"
   WebDAV:
     InternalURLs:
       "http://localhost:9003/": {}
-   ExternalURL: "https://172.17.0.2:9002/"
+   ExternalURL: "https://192.168.1.125:9002/"
   WebDAVDownload:
     InternalURLs:
       "http://localhost:9003/": {}
-   ExternalURL: "https://172.17.0.2:9002/"
+   ExternalURL: "https://192.168.1.125:9002/"
     InternalURLs:
       "http://localhost:9003/": {}
   Composer:
-   ExternalURL: "https://172.17.0.2:4200"
+   ExternalURL: "https://192.168.1.125:4200"
   Controller:
-   ExternalURL: "https://172.17.0.2:8000"
+   ExternalURL: "https://192.168.1.125:8000"
     InternalURLs:
       "http://localhost:8003": {}
   RailsAPI:
@@ -76,7 +76,7 @@
   AutoSetupNewUsersWithRepository: true
   Workbench:
     SecretKeyBase: 19dvt4z7b4v0feeyzo8kzwfq08w97vsefjbllk6qoybvp23dms7bc6sii0bwovs3tgnl4me5mknt52
-   ArvadosDocsite: http://172.17.0.2:8001/
+   ArvadosDocsite: http://192.168.1.125:8001/
   Git:
     GitCommand: /usr/share/gitolite3/gitolite-shell
     GitoliteHome: /var/lib/arvados/git
```

## nginx.conf

```
--- nginx.conf.dev    2019-12-26 20:46:44.625148516 -0300
+++ nginx.conf.publicdev    2019-12-26 20:38:16.612543296 -0300
@@ -41,8 +41,8 @@
server {
    listen *:8000 ssl default_server;
    server_name controller;
-   ssl_certificate "/var/lib/arvados/server-cert-172.17.0.2.pem";
-   ssl_certificate_key "/var/lib/arvados/server-cert-172.17.0.2.key";
+   ssl_certificate "/var/lib/arvados/server-cert-192.168.1.125.pem";
+   ssl_certificate_key "/var/lib/arvados/server-cert-192.168.1.125.key";
    location / {
        proxy_pass http://controller;
        proxy_set_header Host $http_host;
@@ -64,8 +64,8 @@
    proxy_read_timeout    300s;

    ssl    on;
-   ssl_certificate "/var/lib/arvados/server-cert-172.17.0.2.pem";
-   ssl_certificate_key "/var/lib/arvados/server-cert-172.17.0.2.key";
+   ssl_certificate "/var/lib/arvados/server-cert-192.168.1.125.pem";
+   ssl_certificate_key "/var/lib/arvados/server-cert-192.168.1.125.key";

    location / {
        proxy_pass    http://arvados-ws;
@@ -82,8 +82,8 @@
server {
    listen *:3001 ssl default_server;
    server_name workbench2;
-   ssl_certificate "/var/lib/arvados/server-cert-172.17.0.2.pem";
-   ssl_certificate_key "/var/lib/arvados/server-cert-172.17.0.2.key";
+   ssl_certificate "/var/lib/arvados/server-cert-192.168.1.125.pem";
+   ssl_certificate_key "/var/lib/arvados/server-cert-192.168.1.125.key";
    location / {
        proxy_pass http://workbench2;
        proxy_set_header Host $http_host;
@@ -106,8 +106,8 @@
server {
    listen *:9002 ssl default_server;
    server_name keep-web;
-   ssl_certificate "/var/lib/arvados/server-cert-172.17.0.2.pem";
-   ssl_certificate_key "/var/lib/arvados/server-cert-172.17.0.2.key";
+   ssl_certificate "/var/lib/arvados/server-cert-192.168.1.125.pem";
+   ssl_certificate_key "/var/lib/arvados/server-cert-192.168.1.125.key";
    client_max_body_size 0;
    location / {
        proxy_pass http://keep-web;
@@ -124,8 +124,8 @@
server {
    listen *:25101 ssl default_server;
    server_name keepproxy;
-   ssl_certificate "/var/lib/arvados/server-cert-172.17.0.2.pem";
-   ssl_certificate_key "/var/lib/arvados/server-cert-172.17.0.2.key";
+   ssl_certificate "/var/lib/arvados/server-cert-192.168.1.125.pem";
+   ssl_certificate_key "/var/lib/arvados/server-cert-192.168.1.125.key";
    client_max_body_size 128M;
    location / {
        proxy_pass http://keepproxy;
@@ -146,8 +146,8 @@
    proxy_read_timeout 300s;

    ssl on;
-   ssl_certificate "/var/lib/arvados/server-cert-172.17.0.2.pem";
-   ssl_certificate_key "/var/lib/arvados/server-cert-172.17.0.2.key";
+   ssl_certificate "/var/lib/arvados/server-cert-192.168.1.125.pem";
+   ssl_certificate_key "/var/lib/arvados/server-cert-192.168.1.125.key";
    client_max_body_size 50m;

    location / {
```

When trying to upload something using `arv-put` in "publicdev" mode I get an error like the following:

```
[...]
arvados.errors.KeepWriteError: Error writing some blocks: block 9b7fc0e714e89a7c224d80b38d9da9c0+3702 raised K
```

```
eepWriteError (failed to write 9b7fc0e714e89a7c224d80b38d9da9c0 after 2 attempts (wanted 1 copies but wrote 0)
: service https://192.168.1.125:25101/ responded with 413 HTTP/1.1 100 Continue
HTTP/1.1 413 Request Entity Too Large)
```

#### Keepproxy log:

```
2019-12-27_15:06:38.99379 + arv keep_service update --uid xzky2-bi614-8uqa612e0jzqq02 --keep-service '{
2019-12-27_15:06:38.99379 "service_host": "192.168.1.125",
2019-12-27_15:06:38.99379 "service_port": 25101,
2019-12-27_15:06:38.99380 "service_ssl_flag": true,
2019-12-27_15:06:38.99380 "service_type": "proxy"
2019-12-27_15:06:38.99380 }'
2019-12-27_15:07:11.39649 {
2019-12-27_15:07:11.39651 "href": "/keep_services/xzky2-bi614-8uqa612e0jzqq02",
2019-12-27_15:07:11.39652 "kind": "arvados#keepService",
2019-12-27_15:07:11.39652 "etag": "4i4pjh9s9h7eihlg32vny4fum",
2019-12-27_15:07:11.39652 "uuid": "xzky2-bi614-8uqa612e0jzqq02",
2019-12-27_15:07:11.39653 "owner_uuid": "xzky2-tpzed-0000000000000000",
2019-12-27_15:07:11.39653 "created_at": "2019-12-13T15:04:08.866492000Z",
2019-12-27_15:07:11.39653 "modified_by_client_uuid": "xzky2-ozdt8-tfmze53vhjnp8m",
2019-12-27_15:07:11.39655 "modified_by_user_uuid": "xzky2-tpzed-0000000000000000",
2019-12-27_15:07:11.39655 "modified_at": "2019-12-27T15:07:11.054103000Z",
2019-12-27_15:07:11.39655 "service_host": "192.168.1.125",
2019-12-27_15:07:11.39656 "service_port": 25101,
2019-12-27_15:07:11.39656 "service_ssl_flag": true,
2019-12-27_15:07:11.39656 "service_type": "proxy",
2019-12-27_15:07:11.39656 "read_only": false
2019-12-27_15:07:11.39657 }
2019-12-27_15:07:11.43825 + exec /usr/local/bin/keepproxy
2019-12-27_15:07:11.49734 time="2019-12-27T15:07:11Z" level=info msg="keepproxy dev started"
2019-12-27_15:07:18.83813 time="2019-12-27T15:07:18Z" level=info msg="listening at 127.0.0.1:25100"
2019-12-27_15:24:23.18276 time="2019-12-27T15:24:23Z" level=info msg=request RequestID=req-vzwv6r10brkdt6qq800
h remoteAddr="127.0.0.1:52736" reqBytes=3702 reqForwardedFor=192.168.1.107 reqHost="192.168.1.125:25101" reqMe
thod=PUT reqPath=9b7fc0e714e89a7c224d80b38d9da9c0 reqQuery=
2019-12-27_15:24:23.22728 time="2019-12-27T15:24:23Z" level=info msg=request RequestID=req-vzwv6r10brkdt6qq800
h remoteAddr="127.0.0.1:52748" reqBytes=3702 reqForwardedFor=172.17.0.1 reqHost="192.168.1.125:25101" reqMetho
d=PUT reqPath=9b7fc0e714e89a7c224d80b38d9da9c0 reqQuery=
2019-12-27_15:24:23.22732 time="2019-12-27T15:24:23Z" level=info msg="proxy loop detected (request has Via: \"
HTTP/1.0 keepproxy\"): perhaps keepproxy is misidentified by gateway config as an external client, or its keep
_services record does not have service_type=proxy?"
2019-12-27_15:24:23.22736 time="2019-12-27T15:24:23Z" level=info msg=response RequestID=req-vzwv6r10brkdt6qq800
h remoteAddr="127.0.0.1:52748" reqBytes=3702 reqForwardedFor=172.17.0.1 reqHost="192.168.1.125:25101" reqMeth
od=PUT reqPath=9b7fc0e714e89a7c224d80b38d9da9c0 reqQuery= respBody="loop detected\n" respBytes=14 respStatus="
Internal Server Error" respStatusCode=500 timeToStatus=0.000102 timeTotal=0.000107 timeWriteBody=0.000005
2019-12-27_15:24:23.24343 time="2019-12-27T15:24:23Z" level=info msg=request RequestID=req-vzwv6r10brkdt6qq800
h remoteAddr="127.0.0.1:52754" reqBytes=3702 reqForwardedFor=172.17.0.1 reqHost="192.168.1.125:25101" reqMetho
d=PUT reqPath=9b7fc0e714e89a7c224d80b38d9da9c0 reqQuery=
2019-12-27_15:24:23.24345 time="2019-12-27T15:24:23Z" level=info msg="proxy loop detected (request has Via: \"
HTTP/1.0 keepproxy\"): perhaps keepproxy is misidentified by gateway config as an external client, or its keep
_services record does not have service_type=proxy?"
2019-12-27_15:24:23.24351 time="2019-12-27T15:24:23Z" level=info msg=response RequestID=req-vzwv6r10brkdt6qq800
h remoteAddr="127.0.0.1:52754" reqBytes=3702 reqForwardedFor=172.17.0.1 reqHost="192.168.1.125:25101" reqMeth
od=PUT reqPath=9b7fc0e714e89a7c224d80b38d9da9c0 reqQuery= respBody="loop detected\n" respBytes=14 respStatus="
Internal Server Error" respStatusCode=500 timeToStatus=0.000087 timeTotal=0.000092 timeWriteBody=0.000005
2019-12-27_15:24:23.24689 time="2019-12-27T15:24:23Z" level=info msg=request RequestID=req-vzwv6r10brkdt6qq800
h remoteAddr="127.0.0.1:52760" reqBytes=3702 reqForwardedFor=172.17.0.1 reqHost="192.168.1.125:25101" reqMetho
d=PUT reqPath=9b7fc0e714e89a7c224d80b38d9da9c0 reqQuery=
2019-12-27_15:24:23.24691 time="2019-12-27T15:24:23Z" level=info msg="proxy loop detected (request has Via: \"
HTTP/1.0 keepproxy\"): perhaps keepproxy is misidentified by gateway config as an external client, or its keep
_services record does not have service_type=proxy?"
2019-12-27_15:24:23.24698 time="2019-12-27T15:24:23Z" level=info msg=response RequestID=req-vzwv6r10brkdt6qq800
h remoteAddr="127.0.0.1:52760" reqBytes=3702 reqForwardedFor=172.17.0.1 reqHost="192.168.1.125:25101" reqMeth
od=PUT reqPath=9b7fc0e714e89a7c224d80b38d9da9c0 reqQuery= respBody="loop detected\n" respBytes=14 respStatus="
Internal Server Error" respStatusCode=500 timeToStatus=0.000080 timeTotal=0.000085 timeWriteBody=0.000005
2019-12-27_15:24:23.24725 time="2019-12-27T15:24:23Z" level=info msg="192.168.1.107,127.0.0.1:52736 PUT /9b7fc0e714e89a7c224d80b38d9da9c0 413 3702 1 0 Could not write sufficient replicas: [500] "
2019-12-27_15:24:23.24734 time="2019-12-27T15:24:23Z" level=info msg=response RequestID=req-vzwv6r10brkdt6qq800
h remoteAddr="127.0.0.1:52736" reqBytes=3702 reqForwardedFor=192.168.1.107 reqHost="192.168.1.125:25101" reqM
ethod=PUT reqPath=9b7fc0e714e89a7c224d80b38d9da9c0 reqQuery= respBody="Could not write sufficient replicas: [500] \n" respBytes=44 respStatus="Request Entity Too Large" respStatusCode=413 timeToStatus=0.070566 timeTotal=0.070571 timeWriteBody=0.000005
```

#### #3 - 12/27/2019 08:39 PM - Lucas Di Pentima

I believe the keepproxy's accessible API call is not realizing that it is being called from a proxy so the API returns keepproxy's external url and so the

loop is detected, here's a relevant piece of controller's log:

```
[...]
2019-12-27_20:33:40.37761 {"PID":41,"RequestID":"req-n45japx0f51f2kmtvttd","level":"info","msg":"request","remoteAddr":"127.0.0.1:55808","reqBytes":0,"reqForwardedFor":"192.168.1.107","reqHost":"192.168.1.125:8000","reqMethod":"GET","reqPath":"/arvados/v1/keep_services/accessible","reqQuery":{"alt=json","time":"2019-12-27T20:33:40.377524040Z"}}
2019-12-27_20:33:40.39703 {"PID":41,"RequestID":"req-n45japx0f51f2kmtvttd","level":"info","msg":"response","remoteAddr":"127.0.0.1:55808","reqBytes":0,"reqForwardedFor":"192.168.1.107","reqHost":"192.168.1.125:8000","reqMethod":"GET","reqPath":"/arvados/v1/keep_services/accessible","reqQuery":{"alt=json","respBytes":634,"respStatus":"OK","respStatusCode":200,"time":"2019-12-27T20:33:40.396944480Z","timeToStatus":0.019323,"timeTotal":0.019412,"timeWriteBody":0.000089}}
[...]
2019-12-27_20:33:44.38284 {"PID":41,"RequestID":"req-n45japx0f51f2kmtvttd","level":"info","msg":"request","remoteAddr":"127.0.0.1:55862","reqBytes":0,"reqForwardedFor":"192.168.1.107","reqHost":"192.168.1.125:8000","reqMethod":"GET","reqPath":"/arvados/v1/keep_services/accessible","reqQuery":{"alt=json","time":"2019-12-27T20:33:44.382727121Z"}}
2019-12-27_20:33:44.40316 {"PID":41,"RequestID":"req-n45japx0f51f2kmtvttd","level":"info","msg":"response","remoteAddr":"127.0.0.1:55862","reqBytes":0,"reqForwardedFor":"192.168.1.107","reqHost":"192.168.1.125:8000","reqMethod":"GET","reqPath":"/arvados/v1/keep_services/accessible","reqQuery":{"alt=json","respBytes":634,"respStatus":"OK","respStatusCode":200,"time":"2019-12-27T20:33:44.403021658Z","timeToStatus":0.020209,"timeTotal":0.020289,"timeWriteBody":0.000079}}
[...]
```

The first couple of lines correspond to arv-put's first attempt being executed from 192.168.1.107 and I'm supposing keepproxy's discovery request as well... 4 secs later, a 2nd attempt produces another 2 calls. This is confirmed by the RailsAPI's log:

```
[...]
{"method":"GET","path":"/arvados/v1/keep_services/accessible","format":"json","controller":"Arvados::V1::KeepServicesController","action":"accessible","status":200,"duration":7.4,"view":0.36,"db":1.75,"request_id":"req-n45japx0f51f2kmtvttd","client_ipaddr":"127.0.0.1","client_auth":"xzky2-gj3su-5c5947tygjjynzfe","params":{"alt":"json"},"@timestamp":"2019-12-27T20:33:40.395857252Z","@version":"1","message":"[200] GET /arvados/v1/keep_services/accessible (Arvados::V1::KeepServicesController#accessible)"}
{"method":"GET","path":"/arvados/v1/keep_services/accessible","format":"json","controller":"Arvados::V1::KeepServicesController","action":"accessible","status":200,"duration":7.19,"view":0.31,"db":1.64,"request_id":"req-n45japx0f51f2kmtvttd","client_ipaddr":"127.0.0.1","client_auth":"xzky2-gj3su-5c5947tygjjynzfe","params":{"alt":"json"},"@timestamp":"2019-12-27T20:33:44.401976944Z","@version":"1","message":"[200] GET /arvados/v1/keep_services/accessible (Arvados::V1::KeepServicesController#accessible)"}
[...]
```

Just 2 calls were made to RailsAPI, 4 secs apart from each other. Maybe the controller is caching the results and returning an external client's accessible response to the keepproxy?

#### #4 - 12/30/2019 02:35 PM - Tom Clegg

Lucas Di Pentima wrote:

Maybe the controller is caching the results and returning an external client's accessible response to the keepproxy?

controller doesn't do any caching.

I'm guessing that with this config, when keepproxy connects to nginx using the host's IP address (192.168.1.125), the request is forwarded back into the container according to the docker --publish args, so the remote IP seen by nginx is 192.168.1.125, which is classified as external.

#### #5 - 01/02/2020 03:46 PM - Peter Amstutz

- Target version changed from 2020-01-02 Sprint to 2020-01-15 Sprint

#### #6 - 01/15/2020 03:50 PM - Peter Amstutz

- Target version changed from 2020-01-15 Sprint to 2020-01-29 Sprint

#### #7 - 01/24/2020 09:59 PM - Lucas Di Pentima

Adding debug logging to the API's keep\_services controller I was able to see the following:

#### First accessible request headers (from the external client)

```
"HTTP_USER_AGENT"=>"google-api-python-client/1.6.7 (gzip)"
"HTTP_ACCEPT_ENCODING"=>"gzip"
"HTTP_AUTHORIZATION"=>"OAuth2 <token>"
"HTTP_X_FORWARDED_PROTO"=>"https"
"HTTP_X_REQUEST_ID"=>"req-6mhic2pgyfrerspjlhhg"
"HTTP_ACCEPT"=>"application/json"
"HTTP_X_EXTERNAL_CLIENT"=>"1"
"HTTP_HOST"=>"192.168.1.125:8000"
"HTTP_VIA"=>"HTTP/1.0 arvados-controller"
"HTTP_X_FORWARDED_FOR"=>"192.168.1.107,127.0.0.1:38346"
"HTTP_VERSION"=>"HTTP/1.1"
```

## Second accessible request headers (from keepproxy itself)

```
"HTTP_USER_AGENT"=>"Go-http-client/1.1"
"HTTP_AUTHORIZATION"=>"OAuth2 <token>"
"HTTP_X_FORWARDED_PROTO"=>"https"
"HTTP_X_REQUEST_ID"=>"req-6mhic2pgyfrerspjlhhg"
"HTTP_X_EXTERNAL_CLIENT"=>"1"
"HTTP_HOST"=>"192.168.1.125:8000"
"HTTP_VIA"=>"HTTP/1.0 arvados-controller"
"HTTP_X_FORWARDED_FOR"=>"172.17.0.1,127.0.0.1:38354"
"HTTP_VERSION"=>"HTTP/1.1"
```

192.168.1.125 is the "public" IP assigned to arvbox. On the second request we can see that keepproxy's request is being detected as external because for some reason it's origin IP is 172.17.0.1 (docker0 interface's IP on the host) instead of 172.17.0.2 (arvbox container's IP)

I think the fact that keepproxy uses the external address 192.168.1.125 to connect to the controller to ask for the keep services has something to do with this. One fix for arvbox's case would be to add the default gw's IP as local on nginx.conf

### #8 - 01/24/2020 10:28 PM - Lucas Di Pentima

Updates at [6e89a1027](#) - branch 15944-arvbox-publicdev-fix

- Treats arvbox's default gw IP as internal so that Keep works on publicdev mode.
- Turns off nginx's response caching for easier debugging (when looking at api's server logs for accessible requests, sometimes keepproxy's requests didn't appear on the log).

### #9 - 01/28/2020 06:50 PM - Peter Amstutz

lucas: when keepproxy does the 'accessible' request to the API server, it connects to the controller's public IP, that is, it connects to the nginx again and the nginx sees this incoming connection as being from the docker0 interface's IP from the host, so it marks it as external, forwarding the request again to keepproxy, and then keepproxy detects the loop and fails.

peter: I see. when it connects to the "external" IP the origin is the gateway not the docker bridge IP address. that makes sense.

LGTM!

### #10 - 01/28/2020 07:20 PM - Lucas Di Pentima

- Status changed from *In Progress* to *Resolved*