D4Science Infrastructure - Task #6710

virtual machine with apache 2.4.8 (or higher)

Jan 30, 2017 12:19 PM - Ciro Formisano

Status:	Closed	**Start date:**	Jan 30, 2017
Priority:	Normal	**Due date:**	
Assignee:	_InfraScience Systems Engineer	**% Done:**	100%
Category:		**Estimated time:**	0.00 hour
Target version:	Development Environment Enhancements		
Infrastructure:	Development		

Description

In order to test the HTTPS proxy needed for #5422 I need an instance of apache 2.4.8 (on a dedicated VM or not, it is not important) and the possibility to set the HTTP portlets to HTTPS in the test environment.

Related issues:

Blocks D4Science Infrastructure - Task #5422: develop a proxy service to be u...	**Closed**	**Oct 05, 2016**

History

#1 - Jan 30, 2017 12:20 PM - Ciro Formisano

- Blocks Task #5422: develop a proxy service to be used for access http external pages from the portal domain added

#2 - Jan 30, 2017 03:01 PM - Pasquale Pagano

@massimiliano.assante@isti.cnr.it, do we need another server? Can we allocate it on an existing one?

#3 - Jan 30, 2017 03:25 PM - Massimiliano Assante

- Assignee changed from _InfraScience Systems Engineer to Andrea Dell'Amico

@ciro.formisano@eng.it we have no HTTP portlets, the use case (in my understanding) is that some of the portlets deployed in any Infra Gateway (e.g. https://i-marine.d4science.org) needs to retrieve dynamic content from an external service running over HTTP. In practice, any query originating client-side (on the browser) will be prepended (Portlet Developer does the job) by the https proxy first.

@pasquale.pagano@isti.cnr.it I was wondering if an apache implementation is the optimal solution in this case. If we do so we won't have authorisation nor accounting on the node (s), I'm afraid we're actually opening a proxy to the world.

Also, a single instance of the service will introduce a SPOF in the infra. Since the https proxy service is stateless, it would be beneficial to create 2 instances at least and to use a load balancer on top of it.

#4 - Jan 30, 2017 03:36 PM - Ciro Formisano

@massimiliano.assante@isti.cnr.it: according to the requirements I have, Apache is a good solution. In general it opens a proxy from our internal network to the external (from a security point of view I would be worried by the opposite): anyway, if the HTTP dynamic content is limited we could limit the endpoints.
However I am very far from saying that we have the solution: for this reason I would like to test this idea in an environment similar to the destination one with the portlets that currently are problematic (or a subset of them). If there are more requirements that I could have missed, and could make this solution unapplyable please let me know.

#5 - Jan 31, 2017 12:46 PM - Andrea Dell'Amico

- Assignee changed from Andrea Dell'Amico to _InfraScience Systems Engineer

#6 - Feb 01, 2017 12:41 PM - Andrea Dell'Amico

Ciro Formisano wrote:

> @massimiliano.assante@isti.cnr.it: according to the requirements I have, Apache is a good solution. In general it opens a proxy from our internal network to the external (from a security point of view I would be worried by the opposite): anyway, if the HTTP dynamic content is limited we could limit the endpoints.

The proxy will be accessed by the portlets, correct? So we can limit the access using iptables rules and maybe apache ACLs

> However I am very far from saying that we have the solution: for this reason I would like to test this idea in an environment similar to the

> destination one with the portlets that currently are problematic (or a subset of them). If there are more requirements that I could have missed, and could make this solution unapplyable please let me know.

Is the 2.4.8 version mandatory? Ubuntu 14.04 provides apache 2.4.7 and that's our platform right now.

#7 - Feb 01, 2017 01:07 PM - Ciro Formisano

Hi Andrea,
I am not sure that the proxy will be accessed only by the Portlets. As far as I know it is not a "backend" of the Portal, but it should hide some services accessed directly by the client (browser but not only) that are in HTTP and this leads to errors since the browser expects to use HTTPs. Anyway @francesco.mangiacrapa@isti.cnr.it probably can be more detailed than me.

Concerning the version of apache, unluckily I need some features available from 2.4.8: I have a Debian and the version available in the repositories is 2.4.10 (on which I made the tests)

#8 - Feb 01, 2017 01:14 PM - Andrea Dell'Amico

Ciro Formisano wrote:

> Hi Andrea,
> I am not sure that the proxy will be accessed only by the Portlets. As far as I know it is not a "backend" of the Portal, but it should hide some services accessed directly by the client (browser but not only) that are in HTTP and this leads to errors since the browser expects to use HTTPs. Anyway @francesco.mangiacrapa@isti.cnr.it probably can be more detailed than me.

It that's the case @massimiliano.assante@isti.cnr.it is right: this is going to be a huge security hole because we are activating an open proxy and letting it available to everyone. So it's not a viable solution.

> Concerning the version of apache, unluckily I need some features available from 2.4.8: I have a Debian and the version available in the repositories is 2.4.10 (on which I made the tests)

nginx could be an option, possibly?

#9 - Feb 01, 2017 01:30 PM - Ciro Formisano

> It that's the case @massimiliano.assante@isti.cnr.it is right: this is going to be a huge security hole because we are activating an open proxy and letting it available to everyone. So it's not a viable solution.

The problem you are pointing could be related to the solution, not to its implementation. For the moment we have some services accessed in HTTP on behalf of HTTPS: actually I don't know how these services are protected at portal level and if the new solution adds or remove current security issues (if present). I would evaluate also these problems by the test session.

> nginx could be an option, possibly?

Version 2.4.x is OK for everything except for backreferences of locationmatch (https://httpd.apache.org/docs/2.4/mod/core.html#locationmatch), available from 2.4.8.

#10 - Feb 01, 2017 01:45 PM - Andrea Dell'Amico

Ciro Formisano wrote:

> The problem you are pointing could be related to the solution, not to its implementation. For the moment we have some services accessed in HTTP on behalf of HTTPS: actually I don't know how these services are protected at portal level and if the new solution adds or remove current security issues (if present). I would evaluate also these problems by the test session.

So what you're saying is that only a (fixed) list of target URLs will be accepted?

> Version 2.4.x is OK for everything except for backreferences of locationmatch (https://httpd.apache.org/docs/2.4/mod/core.html#locationmatch), available from 2.4.8.

I've found a PPA for apache2, btw: https://launchpad.net/~ondrej/+archive/ubuntu/apache2?field.series_filter=trusty

#11 - Feb 01, 2017 01:57 PM - Ciro Formisano

> So what you're saying is that only a (fixed) list of target URLs will be accepted?

No, I am saying that at this stage we have to test the candidate solution and evaluate functional and not functional problems (if present). Then we will fix them (if possible) otherwise we will find another solution. Your proposal could be OK but I have to test it.

> I've found a PPA for apache2, btw: https://launchpad.net/~ondrej/+archive/ubuntu/apache2?field.series_filter=trusty

It seems that it provides apache 2.4.25. Should be OK

#12 - Feb 01, 2017 03:18 PM - Francesco Mangiacrapa

@andrea.dellamico@isti.cnr.it and @ciro.formisano@eng.it thanks for your replies

A question:
Can we add one or more configuration (security policies) through which the proxy will be accessible: either from LAN or by a list of ip addresses otherwise by gcube token and/or etc..??
I think yes.
So the security is not a problem (at least I hope).
At the moment, the issue is that It is not possible to access to dynamic HTTP content from our portals.. and a solution could be an https to http proxy adding obviously our policies to the security.

#13 - Feb 01, 2017 03:28 PM - Andrea Dell'Amico

Francesco Mangiacrapa wrote:

> @andrea.dellamico@isti.cnr.it and @ciro.formisano@eng.it thanks for your replies
>
> A question:
> Can we add one or more configuration (security policies) through which the proxy will be accessible: either from LAN or by a list of ip addresses otherwise by gcube token and/or etc..??
> I think yes.
> So the security is not a problem (at least I hope).

It's what I was asking in #note-10. network or hosts restrictions are not a problem. I don't know how it could be possible using a gcube token.

#14 - Feb 01, 2017 03:37 PM - Ciro Formisano

@francesco.mangiacrapa@isti.cnr.it thank you for the reply.

I would only avoid, for the moment, to add any specific security policy while we are performing tests. I think that we will need and I also think that it will not be difficult to define and apply them, but if we apply security at this time I fear that we will have problems to clearly understand if the functionalities of the solution are OK.
For this reason I asked @andrea.dellamico@isti.cnr.it to not think about security for the moment (even if we consider the problem).

#15 - Feb 01, 2017 03:43 PM - Francesco Mangiacrapa

Andrea Dell'Amico wrote:

> It's what I was asking in #note-10. network or hosts restrictions are not a problem. I don't know how it could be possible using a gcube token.

You're right @andrea.dellamico@isti.cnr.it. Security based on gcube-token could be very hard (impossible?) to be implemented/integrated via Apache.
I think that network or hosts restrictions are sufficient even because is not need to perform accounting for the request passed via proxy.. :-) Right? :-)

#16 - Feb 01, 2017 07:33 PM - Pasquale Pagano

two short comments:
a) we have a concrete issue and we need a solution. Please don't try to solve all possible problems.
the issue is that It is not possible to access to dynamic HTTP content from our portals running on HTTPS.
The solution we though as to define an https to http proxy.
I am not an expert but the proxy is only for outgoing calls from our network to the web and from our portal only. So rules can be formulated under those considerations.

b) security is an issue even during the test phase. We cannot open to the world risking to compromise the infrastructure we run. So, the testing phase cannot be started without the right security configuration

#17 - Feb 01, 2017 10:06 PM - Massimiliano Assante

Pasquale Pagano wrote:

> two short comments:
> a) we have a concrete issue and we need a solution. Please don't try to solve all possible problems.
> the issue is that It is not possible to access to dynamic HTTP content from our portals running on HTTPS.

> The solution we though as to define an https to http proxy.
> I am not an expert but(alls from our network to the web and from our portal only. So rules can be formulated under those considerations.

Nope, any browser contacts the proxy from anywhere.

I was clearly wrong but even accepting requests from anywhere we know the target of the calls to accept. So the filters can be set on the target and not on the source.

#18 - Feb 02, 2017 01:48 PM - Andrea Dell'Amico

- Status changed from New to In Progress

In the meantime that we discuss the best approach I'm going to activate a VM with the latest apache and a letsencrypt certificate installed but without any apache virtualhost configuration.

VM hostname and IP are going to be portlet-proxy-d-d4s.d4science.org 146.48.122.6

#19 - Feb 02, 2017 03:17 PM - Andrea Dell'Amico

- Status changed from In Progress to Feedback

- % Done changed from 0 to 100

The VM is running. apache 2.4.25 has been installed, ports 80 and 443 are both open to the world.

The letsencrypt tools are installed but a valid certificate cannot be requested right now because we reached the cap of 20 certificates/week. A self signed certificate is present in the same path that will be used by the valid one:

```
/var/lib/acme/live/portlet-proxy-d-d4s.d4science.org/cert
/var/lib/acme/live/portlet-proxy-d-d4s.d4science.org/chain
/var/lib/acme/live/portlet-proxy-d-d4s.d4science.org/fullchain
/var/lib/acme/live/portlet-proxy-d-d4s.d4science.org/privkey
```

#20 - Feb 02, 2017 03:25 PM - Ciro Formisano

Thank you Andrea,
I will work on that machine as soon as @massimiliano.assante@isti.cnr.it , @francesco.mangiacrapa@isti.cnr.it and me will have agreed on the best solution.
Meanwhile I should get access to this machine: may you help me?

#21 - Feb 02, 2017 03:33 PM - Andrea Dell'Amico

Ciro Formisano wrote:

> Thank you Andrea,
> I will work on that machine as soon as @massimiliano.assante@isti.cnr.it , @francesco.mangiacrapa@isti.cnr.it and me will have agreed on the best solution.
> Meanwhile I should get access to this machine: may you help me?

I forgot to say that you should have ssh access as root already. I've found a ssh key of yours on manage.research-infrastructures.eu. If it's not valid anymore you can add a ssh key into your profile data on the portal.

#22 - Feb 02, 2017 04:06 PM - Ciro Formisano

- Status changed from Feedback to Closed

works: thank you.

#23 - Feb 16, 2017 02:59 PM - Ciro Formisano

- Status changed from Closed to In Progress

- % Done changed from 100 to 50

@francesco.mangiacrapa@isti.cnr.it , please add the details of the new deployment

#24 - Feb 16, 2017 03:05 PM - Francesco Mangiacrapa

@andrea.dellamico@isti.cnr.it ASAP we need a smartgears machine running in the scope "/gcube/devNext/NextNext"
which is the scope where is up a Proxy Portlet for testing it https://dev4.d4science.org/group/nextnext/test-proxy, thx

#25 - Feb 16, 2017 03:12 PM - Andrea Dell'Amico

The VM is going to be converted to smartgears, apache removed

#26 - Feb 16, 2017 03:14 PM - Andrea Dell'Amico

Note that as it's a development server, it will run Java 8.

#27 - Feb 16, 2017 03:27 PM - Andrea Dell'Amico

- Status changed from In Progress to Feedback

The VM is ready. I increased the memory to 2.5GB to accomodate tomcat.
nginx is configured to redirect to tomcat the /httpproxy context.

#28 - Feb 16, 2017 04:51 PM - Ciro Formisano

looking at ghn.log I see several exceptions, continuously produced such as:

```
 ERROR ProfileManager: cannot publish container (see details)
com.sun.xml.internal.ws.client.ClientTransportException: HTTP transport error: java.net.ConnectException: Conn
ection refused (Connection refused)
        at com.sun.xml.internal.ws.transport.http.client.HttpClientTransport.getOutput(HttpClientTransport.jav
a:117) ~[na:1.8.0_121]
        at com.sun.xml.internal.ws.transport.http.client.HttpTransportPipe.process(HttpTransportPipe.java:208)
 ~[na:1.8.0_121]
        at com.sun.xml.internal.ws.transport.http.client.HttpTransportPipe.processRequest(HttpTransportPipe.ja
va:130) ~[na:1.8.0_121]
        at com.sun.xml.internal.ws.transport.DeferredTransportPipe.processRequest(DeferredTransportPipe.java:1
24) ~[na:1.8.0_121]
        at com.sun.xml.internal.ws.api.pipe.Fiber.__doRun(Fiber.java:1121) ~[na:1.8.0_121]
        at com.sun.xml.internal.ws.api.pipe.Fiber._doRun(Fiber.java:1035) ~[na:1.8.0_121]
        at com.sun.xml.internal.ws.api.pipe.Fiber.doRun(Fiber.java:1004) ~[na:1.8.0_121]
        at com.sun.xml.internal.ws.api.pipe.Fiber.runSync(Fiber.java:862) ~[na:1.8.0_121]
        at com.sun.xml.internal.ws.client.Stub.process(Stub.java:448) ~[na:1.8.0_121]
        at com.sun.xml.internal.ws.client.sei.SEIStub.doProcess(SEIStub.java:178) ~[na:1.8.0_121]
        at com.sun.xml.internal.ws.client.sei.SyncMethodHandler.invoke(SyncMethodHandler.java:93) ~[na:1.8.0_1
21]
        at com.sun.xml.internal.ws.client.sei.SyncMethodHandler.invoke(SyncMethodHandler.java:77) ~[na:1.8.0_1
21]
        at com.sun.xml.internal.ws.client.sei.SEIStub.invoke(SEIStub.java:147) ~[na:1.8.0_121]
        at com.sun.proxy.$Proxy48.update(Unknown Source) ~[na:na]
        at sun.reflect.GeneratedMethodAccessor54.invoke(Unknown Source) ~[na:na]
```

#29 - Feb 16, 2017 07:26 PM - Tommaso Piccioli

Tomcat autodeploy was missed, could you please check again?

#30 - Feb 17, 2017 05:29 PM - Ciro Formisano

- Status changed from Feedback to Closed

- % Done changed from 50 to 100

Works. Thank you!