

Infrastructure - Task #3419

CNRead.describe() does not return Content-Length header

2012-12-12 16:33 - Roger Dahl

Status:	In Progress	Start date:	2012-12-12
Priority:	Normal	Due date:	
Assignee:	Chris Jones	% Done:	30%
Category:	Metacat	Estimated time:	0.00 hour
Target version:		Story Points:	
Milestone:	None		
Product Version:	*		

Description
dahl@vm-dataone:~/d1/d1_python/d1_client_onedrive/src\$ curl -v -X HEAD
https://cn-stage.test.dataone.org/cn/v1/object/SEF001_024MTBD004R00_20060719.50.5
...
< HTTP/1.1 200 OK
< Date: Wed, 12 Dec 2012 16:19:10 GMT
< Server: Apache/2.2.14 (Ubuntu)
< DataONE-Checksum: MD5,501f259043b9cbfdfaa4ed3944f93698
< Last-Modified: Thu, 01 Jan 1970 00:00:00 GMT
< DataONE-ObjectFormat: eml://ecoinformatics.org/eml-2.0.1
< DataONE-SerialVersion: 1
< Vary: Accept-Encoding
< Content-Type: text/xml;charset=UTF-8
* no chunk, no close, no size. Assume close to signal end

History

#1 - 2012-12-12 16:38 - Roger Dahl

- Description updated

#2 - 2012-12-12 18:20 - Ben Leinfelder

When I do this locally, I get more headers:

```
leinfelder$ curl -v -X HEAD http://localhost:8080/knb/d1/cn/v1/object/jscientist.7.34
* About to connect() to localhost port 8080 (#0)
* Trying ::1... connected
* Connected to localhost (::1) port 8080 (#0)

HEAD /knb/d1/cn/v1/object/jscientist.7.34 HTTP/1.1
User-Agent: curl/7.21.4 (universal-apple-darwin11.0) libcurl/7.21.4 OpenSSL/0.9.8r zlib/1.2.5
Host: localhost:8080
Accept: /

< HTTP/1.1 200 OK
< Server: Apache-Coyote/1.1
< DataONE-Checksum: MD5,0c8061ad0e1f4976a1165359a37a6100
< Last-Modified: 2012-06-18T22:54:53.133+00:00
< DataONE-ObjectFormat: eml://ecoinformatics.org/eml-2.1.0
< DataONE-SerialVersion: 1
< Content-Type: text/xml
< Content-Length: 7359
< Date: Wed, 12 Dec 2012 18:15:59 GMT
<
* transfer closed with 7359 bytes remaining to read
* Closing connection #0
curl: (18) transfer closed with 7359 bytes remaining to read
```

#3 - 2012-12-12 18:23 - Ben Leinfelder

Note I am going directly to Tomcat locally and not through Apache/mod_jk.

#4 - 2013-02-21 18:41 - Chris Jones

- Status changed from New to In Progress

Trying this against a Metacat Member Node, it works fine:

```
curl -v -X HEAD -k -o - - "https://knb.ecoinformatics.org/knb/d1/mn/v1/object/Akasha.9.1"
```

...

```
HEAD /knb/d1/mn/v1/object/Akasha.9.1 HTTP/1.1
User-Agent: curl/7.21.4 (universal-apple-darwin11.0) libcurl/7.21.4 OpenSSL/0.9.8r zlib/1.2.5
Host: knb.ecoinformatics.org
Accept: /
```

```
< HTTP/1.1 200 OK
< Date: Thu, 21 Feb 2013 18:31:16 GMT
< Server: Apache/2.2.14 (Ubuntu)
< DataONE-Checksum: MD5,2044891fca681f00a45098a360f22ab7
< Last-Modified: Thu, 01 Jan 1970 00:00:00 GMT
< DataONE-ObjectFormat: text/csv
< DataONE-SerialVersion: 1
< Content-Length: 1337
< Vary: Accept-Encoding
< Content-Type: text/xml
<
* SSLv3, TLS alert, Client hello (1):
* transfer closed with 1337 bytes remaining to read
* Closing connection #0
* SSLv3, TLS alert, Client hello (1):
```

Trying against the CN, we get the error:

```
curl -v -X HEAD -k -o - - "https://cn-ucsb-1.dataone.org/cn/v1/object/Akasha.9.1"
```

...

```
HEAD /cn/v1/object/Akasha.9.1 HTTP/1.1
User-Agent: curl/7.21.4 (universal-apple-darwin11.0) libcurl/7.21.4 OpenSSL/0.9.8r zlib/1.2.5
Host: cn-ucsb-1.dataone.org
Accept: /
```

```
< HTTP/1.1 200 OK
< Date: Thu, 21 Feb 2013 18:33:54 GMT
< Server: Apache/2.2.14 (Ubuntu)
< DataONE-Checksum: MD5,2044891fca681f00a45098a360f22ab7
< Last-Modified: Thu, 01 Jan 1970 00:00:00 GMT
< DataONE-ObjectFormat: text/csv
< DataONE-SerialVersion: 1
< Vary: Accept-Encoding
< Content-Type: text/xml; charset=UTF-8
* no chunk, no close, no size. Assume close to signal end
<
* SSLv3, TLS alert, Client hello (1):
* Closing connection #0
* SSLv3, TLS alert, Client hello (1):
```

So, we need to look at this from the CN proxy perspective. We know that Metacat is producing the correct HTTP header since this is shared code between the MN and CN (D1NodeService), and so we need to do the same test where we enable http.wire logging on the CN to see where we might be losing the header.