

TURL lifetime in CASTOR

I report here a clear explanation of this question by Flavia:

Problem raised by a user: copying a file to CASTOR using the dCache client command 'srmcp' gives error:

```
...
GridftpClient: Was not able to send checksum
value:org.globus.ftp.exception.ServerException: Server refused performing the
request. Custom message: (error code 1) [Nested exception message: Custom
message: Unexpected reply: 500 Invalid command.] [Nested exception is
org.globus.ftp.exception.UnexpectedReplyCodeException: Custom message:
Unexpected reply: 500 Invalid command.]
GridftpClient: waiting for completion of transfer
...
```

Flavia: this is the well-known problem with the dCache srmcp client and CASTOR. When CASTOR is configured with the so-called internal gridftp, the TURL returned by the SRM server is only valid within one gridftp session. srmcp requires at least 2 gridftp sessions to make a transfer since the first session is used to verify the checksum. After the first gridftp session the TURL is made invalid by CASTOR making the transfer fail. I guess the SRM server you are using is configured to use the "internal" gridftp. You should ask for the server to be configured for "external" gridftp. In such a case, TURLs are always valid. They do not expire after the first gridftp session is closed.

User: Will it affect anything else if I request this change? (for example FTS?)

No. It should not break anything. Both GFAL/lcg-utils and FTS are able to operate with both internal and external gridftp since they have been already modified to do all their business within one gridftp session (therefore you have no problems with those clients at the moment). The dCache developers refused to change their code (since this implied a change also on the server side to correctly implement the srmCopy request) with the justification that CASTOR was "abusing" the SRM specs. In fact, following the specs, a TURL MUST be valid for the requested pin time and cannot expire before.

Simone: CASTOR at CERN was changed to use the internal gridftp few months ago. Why this was done? Performance?

Flavia: It is not for performance reasons as far as I know. It is for a better internal management. But I understood that the CASTOR team is planning to go back to "external gridftp" since they are having more trouble than advantages. They will do otherwise for the internal business they need to keep under control.

User: he saw this error also with Bestman client.

Timur corrected what Flavia said before: The srm-cp command only makes a single connection to Castor. The GridFTP server advertises a checksumming ability. The srm-cp comment attempts to utilize this but the functionality seems problematic and fails. Subsequent to the checksum failing, srm-cp attempts to continue only to discover the TURL is now invalid.

Flavia asked if it were possible to avoid the checksum transfer stage? Timur said it was; the server should refrain from publishing its support for the checksum extension: the checksum stage is only undertaken if the server advertises its support.

Flavia would take this issue back to the Castor team for further investigation; it seems that the Bestman client also shows the same problems and there was a similar issue with StoRM.

More on this topic: Marteen asks: is it not possible for CASTOR to allow the internal TURL to be used 2 or 3 times? I.e. implement a counter? And Andrea asks: This makes me wonder what happens if the TURL is not

PreStagingTestsCastorTurlLifetime < LCG < TWiki

used at all. That is, for example if I issue an srmPrepareToGet and nothing else after. Would the TURL linger forever?

Olof: No, I think it times out after 1min or so but devvers can confirm. The internal gsiftp TURL is associated with a disk server transfer slot so I don't think it will wait for long.

Andrea asks: why the external GridFTP is much heavier on the disk server than the internal GridFTP? Is the memory usage per transfer different? Or the internal one is just a way to limit the number of concurrent transfers by queuing those which cannot be processed at the moment?

-- ElisaLanciotti - 03 Mar 2009

This topic: LCG > PreStagingTestsCastorTurlLifetime

Topic revision: r1 - 2009-06-05 - AndreaSciaba



Copyright &© 2008-2019 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Send feedback