

globus-url-copy from dcache to non-dcache balanced sites

Problem reported and workaround provided by **Vincenzo Vagnoni** (CNAF/INFN)

In the past we noticed problems in data traffic from dcache sites to CNAF, either not working at all or bandwidth problems. I think that this is due to a java feature which I try to briefly explain, and maybe some site admins are not aware of.

The java VM by default performs DNS lookups caching forever. I.e., if for example one issues a globus-url-copy or similar from a dcache site against a DNS balanced set of gridftp servers, java (and hence dcache) performs just one lookup and caches it, and will use it forever as IP address (keeping it in memory, hence this will hold until the next restart of dcache pools). Then, what could happen if e.g. transferring from dcache sites to CERN castor frontend pools, or CNAF, or whatever is based on a balanced or round-robin or random DNS answer will not work as expected. Mainly two situations:

1. a given dcache instance will send data always to just one gridftp server, always the same as cached by the java VM the first time it queried the DNS. In this case this might result in reduced throughput.
2. if the cached IP address in the meanwhile changed, or the machine was excluded from the pool for hardware problems or dismission or whatever, it will not work at all, since the address will no longer be available on the target (non-dcache) site.

This might be a valid explanation and it may be worth to explore it. Of course, dcache can override this feature (which is wanted by java to avoid IP spoofing but screws up DNS alises mechanisms) internally. But if it does not do it, the default java behaviour is kept.

This default behaviour can be changed by altering an option in the java.security configuration file (paying attention to modify the correct one, since different java versions might be installed on a given machine) on the dcache machines. The option to be modified is `networkaddress.cache.ttl=0`. By default it is set to `-1`, which means infinite caching.

-- Flavia Donno - 20 Jun 2007

This topic: LCG > GSSDDCACHEDNS

Topic revision: r1 - 2007-06-20 - FlaviaDonno



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

or Ideas, requests, problems regarding TWiki? use Discourse or Send feedback