

rCUDA Evaluation Final Report

27 August 2014

Version 1.0

The remote CUDA middleware (rCUDA) is a library providing a set of wrappers around the standard CUDA calls that allows them to be redirected to a remote GPU, hence providing a virtual GPU on a virtual machine or on hardware with no GPU attached. An evaluation of version 4.1 of the library was performed to determine if the library successfully and efficiently provides remote GPU computation. The evaluation was by the Technology Investigation Service (TIS) evaluation team led by George Butler. The other members of the evaluation team were Peter Enstrom, Tony Cole, and Jason Charcalla.

The team chose to evaluate rCUDA due to the growing importance of virtualization within XSEDE. Although several other platforms for GPU virtualization exist (DS-CUDA, GVirtuS), rCUDA was chosen first for evaluation. The evaluation team discovered that rCUDA provided an effective platform for remote GPU computation as well as providing GPU computation within a VM. However, some error conditions were not handled as cleanly as would have been desired. Specifically, connections dropped due to firewalls were not always reported clearly, and could potentially result in the middleware needing to be manually restarted. Finally, it should be noted that the process for acquiring rCUDA is somewhat cumbersome, in that the developers must manually approve every request for the middleware. Despite these minor inconveniences the team recommends the use of rCUDA within XSEDE.

Team recommends middleware for remote GPU computation: YES