

Google Summer Of Code 2014

Ganglia GPU Monitoring Enhancement Project Project Report With Change Details

Student:

Md Ali Ahsan Rana

Mentor:

Rajat Phull

Ganglia Admins:

Nick Satterly, Bernard Lee, Daniel Pocock

Summary:

- NVML v340.25 driver based metrics support.
- CUDA 6.5 based metric support. For incompatible CUDA installation, those metric won't be shown up on ganglia-web.
- Number Of Device(GPU) generalization for custom php graph support on ganglia-web module. Which doesn't require any patch.
- Visualization improvement by adding max limit lines on certain metrics to show better max/usage ratio to users.
- Modifications affect 'nvidia gpu' module and 'ganglia-web' core module.

Installation:

- Same as current instruction.
https://github.com/ranacseruet/gmond_python_modules/tree/master/gpu/nvidia .
However, the last step, 'Patch Ganglia Web' isn't required anymore.

Modifications Details:

Metrics Added:

- '_ecc_db_error' : Double Bit ECC Error.
- '_ecc_sb_error' : Single Bit ECC Error.
- '_power_violation_report' : GPU Power violation.
- '_encoder_util' : Encoder Utilization.
- '_decoder_util' : Decoder Utilization.
- '_bar1_memory' : Bar1 Memory.
- '_shutdown_temp' : Shutdown Temperature.
- '_slowdown_temp': Slowdown Temperature.

Metrics Modified:

- '_graphics_speed' to 'graphics_clock_report'. Reason: Changed 'speed' to 'clock' for better meaningful name. Added '_report' for custom php graph file support.
- '_sm_speed' to '_sm_clock_report'. Reason: same as above.
- '_mem_speed' to '_mem_clock_report'. Reason: same as above.

- ‘_mem_used’ to ‘_fb_memory’. Reason: match with NVML and eliminate ambiguity between ‘Frame Buffer Memory’ and ‘Bar1 Memory’.
- ‘_power_usage’ to ‘_power_usage_report’. Reason: custom php graph file support.
- ‘_max_graphics_speed’ to ‘_max_graphics_clock’. Reason: Changed ‘speed’ to ‘clock’ for better meaningful name.
- ‘_max_sm_speed’ to ‘_max_sm_clock’. Reason: as above.
- ‘_max_mem_speed’ to ‘_max_mem_clock’. Reason: as above.
- Moved ‘_max*’ metrics from graph to top constant metrics section(host overview).

Metrics Deleted:

- power_man_mode.
- perf_state.

Custom Graph Modifications:

- Clear old php graph files which wasn’t working.
- Generalization of No# Of GPUs for which custom graphs can be supported. Now, single php file can render custom graph for each GPU#.
- Added Max line on some graphs to provide better overview of limit/usage ratio.

Feedback Required:

Generalization of No# Of GPUs:

Before, for cluster node with more than one gpus, custom graph for a metric(say ‘util’) had to be named as gpu0_util.php, gpu1_util.php, gpu2_util.php , which was limited to the hard coded numbers added up to. Thus, to make this process simpler, I have added an small modification under ‘ganglia-web’ module, which shouldn’t effect any additional functionality, but facilitate this particular issue to all other modules(which could have multiple device attached in single node) as well. It will be very helpful if admins can review it and provide feedbacks or suggestions if there is some better way to handle such scenario.

Capture Useful GPU Events:

Not added to Module yet. Worked on capturing GPU events which are triggered in case of certain events(such as when power violation occurs or ecc error occurs). However, we didn't find a suitable way to plot them on graph. As we would need a way to plot event details on the graph as well. We already sent an email to community, but didn't get any suggestion so far. However, if admins can suggest some proper way, we sure can bring this feature back in soon. This is what we will check for the next couple of days as well.

Supporting Tasks:

- Created automated task runner via cron job so that some demo computations run to generation usage data on graphs.
- Test case for ECC errors - generate ecc errors manually to populate and verify graph data.
- Deployed and tested implementation on two separate 2-GPU based cluster.