WORKSHOP: OPTIMIZE YOUR PMDK-BASED APPLICATION USING INTEL® VTUNE™ PROFILER

Jennifer DiMatteo
Kevin O'Leary
Technical Consulting Engineers
Intel Corporation

Storage Performance Development Kit (SPDK)
Persistent Memory Development Kit (PMDK)
Intel® VTune™ Profiler

Virtual Forum

AGENDA

01 **Sample: matrix**
Demonstration of analyzing the matrix sample with Intel® VTune™ Profiler

02 **Sample: wordcount with mapreduce**
Demonstration of analyzing the PMDK workcount sample

03 **Additional Intel® VTune™ Profiler Demonstration**
Further demonstration of usage and Q&A

SAMPLE: MATRIX

Tried and true, this sample tells a story about using Intel VTune Profiler to address PMDK issues.

https://software.intel.com/content/www/us/en/develop/documentation/vtune-cookbook/top/tuning-recipes/pmdk-application-overhead.html

SAMPLE: WORDCOUNT

Wordcount with mapreduce sample available in the pmdk-examples github repository:

https://github.com/pmem/pmdk-examples/tree/master/mapreduce

The data used for this demo is a large number of text files.

RUNNING THE SAMPLE

After building the mapreduce sample, follow these steps to run a collection with the latest version of Intel® VTune™ Profiler:

1. Create a folder containing the text used for the word count. In this example, 200 text files each around 1MB were used.
2. Create a script to run the following commands:
    1. $ ./wordcount /mnt/pmem1/../wcfile **load** -d=/local/wcfiles/
        1. Where wcfile is the pmem object created and /local/wcfiles is the folder created in step 1
    2. $ ./wordcount /mnt/pmem1/../wcfile **run** -nm=96 -nr=96
        1. Where –nm and –nr are the number of threads for map and reduce, respectively
3. source /opt/intel/vtune_profiler/vtune-vars.sh
4. Launch the script with VTune Profiler:
    1. $ vtune –collect memory-access -- ./wordcount_script.sh

Demo

Review the pre-collected results

Storage Performance Development Kit (SPDK)
Persistent Memory Development Kit (PMDK)
Intel® VTune™ Profiler

Virtual Forum