```
cpu cores       : 68
apicid          : 0
initial apicid  : 0
fpu             : yes
fpu_exception   : yes
cpuid level     : 13
wp              : yes
flags           : fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pat
pse36 clflush dts acpi mmx fxsr sse sse2 ss ht tm pbe syscall nx pdpe1gb rdtscp lm
constant_tsc arch_perfmon pebs bts rep_good nopl xtopology nonstop_tsc aperfmperf
eagerfpu pni pclmulqdq dtes64 monitor ds_cpl est tm2 ssse3 fma cx16 xtpr pdcm sse4_1
sse4_2 x2apic movbe popcnt tsc_deadline_timer aes xsave avx f16c rdrand lahf_lm abm
3dnowprefetch ring3mwait epb ibrs ibpb fsgsbase tsc_adjust bmi1 avx2 smep bmi2 erms
avx512f rdseed adx avx512pf avx512er avx512cd xsaveopt dtherm ida arat pln pts
spec_ctrl
bogomips        : 2799.83
clflush size    : 64
cache_alignment : 64
address sizes   : 46 bits physical, 48 bits virtual
power management:

processor       : 1
vendor_id       : GenuineIntel
cpu family      : 6
model           : 87
model name      : Intel(R) Xeon Phi(TM) CPU 7250 @ 1.40GHz
stepping        : 1
microcode       : 0x1b6
cpu MHz         : 1176.806
cache size      : 1024 KB
physical id     : 0
siblings        : 272
core id         : 1
cpu cores       : 68
apicid          : 4
initial apicid  : 4
fpu             : yes
fpu_exception   : yes
cpuid level     : 13
wp              : yes
flags           : fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pat
pse36 clflush dts acpi mmx fxsr sse sse2 ss ht tm pbe syscall nx pdpe1gb rdtscp lm
constant_tsc arch_perfmon pebs bts rep_good nopl xtopology nonstop_tsc aperfmperf
eagerfpu pni pclmulqdq dtes64 monitor ds_cpl est tm2 ssse3 fma cx16 xtpr pdcm sse4_1
sse4_2 x2apic movbe popcnt tsc_deadline_timer aes xsave avx f16c rdrand lahf_lm abm
3dnowprefetch ring3mwait epb ibrs ibpb fsgsbase tsc_adjust bmi1 avx2 smep bmi2 erms
avx512f rdseed adx avx512pf avx512er avx512cd xsaveopt dtherm ida arat pln pts
spec_ctrl
bogomips        : 2799.83
clflush size    : 64
cache_alignment : 64
address sizes   : 46 bits physical, 48 bits virtual
power management:
```

```
import os
os.environ["MKL_NUM_THREADS"]="1"
os.environ["KMP_SETTINGS"]="1"
#os.environ["KMP_PLACE_THREADS"]="60C,3T"
os.environ["KMP_AFFINITY"]="granularity=fine,compact,1,0"
os.environ["KMP_DUPLICATE_LIB_OK"]="TRUE"
os.environ["KMP_BLOCKTIME"]="3"
os.environ["OMP_NUM_THREADS"]="136"
import tensorflow as tf
session_conf = tf.ConfigProto(intra_op_parallelism_threads=136,
inter_op_parallelism_threads=1)
```

VTUNE APS

Your application might underutilize the available logical CPU cores

because of insufficient parallel work, blocking on synchronization, or too much I/O. Perform function or source line-level profiling with tools like Intel® VTune™ Amplifier to discover why the CPU is underutilized.

	Current run	Target	Delta
CPU Utilization	1.70%	>90%	
Back-End Stalls	40.50%	<20%	
SIMD Instr. per Cycle	0.02	>1	

Application: *python3*
Report creation date: *2018-09-11 13:11:11*
OpenMP threads: *8*
HW Platform: *Intel(R) Processor code named Knights Landing*
Logical Core Count per node: *272*
Collector type: *Event-based counting driver*

577.31s
Elapsed Time

3.75
CPI

CPU Utilization
1.70%

Average CPU Utilization
4.65 Out of 272.00 logical CPUs

Back-End Stalls
40.50% of pipeline slots

L2 Hit Bound
5.40% of cycles

L2 Miss Bound
28.90% of cycles

Average DRAM Bandwidth
1.65 GB/s

Average MCDRAM Bandwidth

SIMD Instr. per Cycle
0.02

FP Instruction Mix
% of Packed SIMD Instr.: 95.40%
% of Scalar SIMD Instr.: 4.60%

Memory Footprint
Resident total: 594.67 MB
Virtual total: 34083.09 MB

```
| Your application might underutilize the available logical CPU cores
| because of insufficient parallel work, blocking on synchronization, or too much I/
O. Perform function or source line-level profiling with tools like Intel(R) VTune(TM)
Amplifier to discover why the CPU is underutilized.
|
  Elapsed time:              577.31 sec
  CPI Rate:                    3.75
| The CPI value may be too high.
| This could be caused by such issues as memory stalls, instruction starvation,
| branch misprediction, or long latency instructions.
| Use Intel(R) VTune(TM) Amplifier General Exploration analysis to specify
| particular reasons of high CPI.
  CPU Utilization:                                  1.70%
| Your application might underutilize the available logical CPU cores because of
| insufficient parallel work, blocking on synchronization, or too much I/O.
| Perform function or source line-level profiling with tools like Intel(R)
| VTune(TM) Amplifier to discover why the CPU is underutilized.
    Average CPU Utilization:                        4.65   Out of  272 logical CPUs
  Back-End Stalls:                                 40.50%
| A significant proportion of pipeline slots remain empty. When operations take
| too long in the back-end, they introduce bubbles in the pipeline that
| ultimately cause fewer pipeline slots containing useful work to be retired per
| cycle than the machine is capable of supporting. This opportunity cost results
| in slower execution. Long-latency operations like division and memory
| operations can cause this, as can too many operations being directed to a
| single execution port (for example, more multiplication operations arriving in
| the back-end per cycle than the execution unit can support). Explore second
| level metrics or use Intel(R) VTune(TM) Amplifier Memory Access analysis to
| learn more.
    L2 Hit Bound:                                  5.40% of cycles
    L2 Miss Bound:                                28.90% of cycles
| A significant proportion of cycles is being spent waiting for L2 load misses
| to be serviced. Possible optimizations are to reduce data working set size,
| improve data access locality, blocking and consuming data in chunks that fit
| in the L2, or better exploit hardware prefetchers. Use Intel(R) VTune(TM)
| Amplifier XE Memory Access analysis to learn more on possible reasons and next
| steps on optimization.
    Average DRAM Bandwidth:                        1.65  GB/s
    Average MCDRAM Bandwidth:                      0.01  GB/s
  SIMD Instructions per Cycle:                       0.02
| The metric value indicates that FPU might be underutilized. This can be a
| result of significant fraction of non-floating point instructions, inefficient
| vectorization because of legacy vector instruction set or memory access
| pattern issues, or different kinds of stalls in the code execution. Explore
| second level metrics to identify the next steps in FPU usage improvements.
       % of Packed SIMD Instr.:                     95.40%
       % of Scalar SIMD Instr.:                      4.60%
 Memory Footprint:
  Resident:        594.67 MB
  Virtual:       34083.09 MB

Graphical representation of this data is available in the HTML report: /home/potero/
Downloads/JN2/aps_report_20180911_132144.html
```

```
KMP_TASK_STEALING_CONSTRAINT=1
KMP_TEAMS_THREAD_LIMIT=272
KMP_TOPOLOGY_METHOD=all
KMP_USER_LEVEL_MWAIT=false
KMP_VERSION=false
KMP_WARNINGS=true
OMP_ALLOCATOR=omp_default_mem_alloc
OMP_CANCELLATION=false
OMP_DEFAULT_DEVICE=0
OMP_DISPLAY_ENV=false
OMP_DYNAMIC=false
OMP_MAX_ACTIVE_LEVELS=2147483647
OMP_MAX_TASK_PRIORITY=0
OMP_NESTED=false
OMP_NUM_THREADS='136'
OMP_PLACES: value is not defined
OMP_PROC_BIND='intel'
OMP_SCHEDULE='static'
OMP_STACKSIZE=4M
OMP_THREAD_LIMIT=2147483647
OMP_TOOL=enabled
OMP_TOOL_LIBRARIES: value is not defined
OMP_WAIT_POLICY=PASSIVE
KMP_AFFINITY='noverbose,warnings,respect,granularity=fine,compact,1,0'
```