Contribution ID: **39** Type: **not specified**

Cache Aware Scheduling

Scheduler Micro conference Proposal

Title:
Cache Aware Scheduling

Presenters:
Tim Chen (tim.c.chen@linux.intel.com)
Chen Yu (yu.c.chen@intel.com)

We have proposed RFC patch series that implemented cache aware scheduling.
The primary motivation is to keep threads sharing data together in the
same last level cache domain, to reduce cache bouncing.

We'll like community feedback on issues that we need to address to take
it upstream.

Discussion Areas:
1. Overview of the use cases that motivate this feature, and current performance numbers.
2. Whether the basic approach of current patches look good to everyone.
3. Whether this feature should be extended to aggregate threads in a single process, but also
to processes communicating via pipes/sockets or sharing memory.
4. Will it be a good idea to use the memory scanning mechanisms in NUMA balancing
to get an ideas on how many tasks are sharing data (from numa_group)? And
perhaps also estimate the extent of shared data?
5. Whether the load aggregation policy implemented in the current patch series can be improved.

Key Participants:
Peter Zijlstra (Intel)
Len Brown (Intel)
Prateek Nayak (AMD)
Madadi Vineeth Reddy (IBM)
Vincent Guittot (Linaro)

Primary author: Mr CHEN, Tim (Intel)

Co-author: Mr CHEN, Yu (Intel)

Presenter: Mr CHEN, Tim (Intel)

Session Classification: Scheduler and Real-Time MC

Track Classification: Scheduler and Real-Time MC