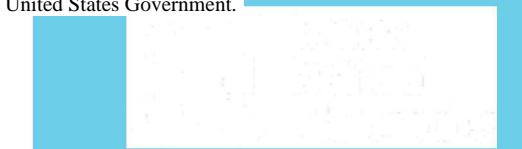


# User Level Threading: A Two-Hat Perspective



*PRESENTED BY*

Stephen Olivier

SNL R&A # (SAND, PR):



UNLIMITED RELEASE

UNCLASSIFIED



Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc. for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

## Hat #1: ULT Implementor (Qthreads)

Fundamental idea: Problem-centric not machine-centric computation

- Don't try to program each application directly to each machine's node topology
- Instead, decompose programs into computational tasks
- Runtime system maps the tasks onto the available execution resources

Additional value-add: full-empty bits for flexible synchronization

- Tera MTA / Cray XMT architectures provided them in hardware
- Qthreads implements in software for similar use cases, e.g., producer-consumer

Key interactions

- On-node tasking layer for Cray's Chapel language
- Collaboration with ANL on {Qthreads, Argobots}+{MPICH, OpenMPI}

<https://github.com/qthreads/>

## Hat #2: ULT Customer (OpenMP Tasking Chair)

Part of the OpenMP\* API spec. since OpenMP 3.0 (2008!)

- Complements original OpenMP loop-level parallelism
- Tasks are executed by OpenMP threads in a team
- Includes OpenMP device (accelerator) offload operations

From syntax and semantics to implementation

- Task support in vendor/LLVM runtimes looks a lot like ULT
- Alternative ULT runtimes (Nanos/OmpSs, Argobots/BOLT)

Continuing challenge

- Cycle of “User adoption” ←/→ “Quality implementations”

