THE UNIVERSITY OF TENNESSEE KNOXVILLE

# PAPI

## THE PERFORMANCE APPLICATION PROGRAMMING INTERFACE

The Performance Application Programming Interface (PAPI) provides a consistent interface (and methodology) for hardware performance counters found across a compute system: i. e., CPUs, GPUs, on- and off-chip memory, interconnects, I/O system, file system, and energy/power. PAPI enables software engineers to see, in near real time, the relationship between software performance and hardware events across the entire compute system.


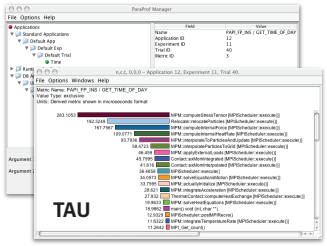


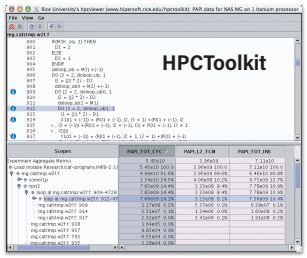


**MORE**

CrayPAT
ompP Tools
Open|SpeedShop
PaRSEC
Scalasca
SCORE-P

### PERFORMANCE ANALYSIS TOOLS

## STANDARD FEATURES

- Standardized Performance Metrics
- Easy Access to Platform-Specific Metrics
- Multiplexed Event Measurement
- Dispatch on Overflow
- Overflow & Profiling on Multiple Simultaneous Events
- Bindings for C, Fortran and Matlab
- User Definable Metrics derived from Platform-Specific Metrics
- Support for Virtual Computing Environments
- Performance Counter Monitoring at Task Granularity for the PaRSEC Dataflow Runtime
- Software-defined Event (SDE) support

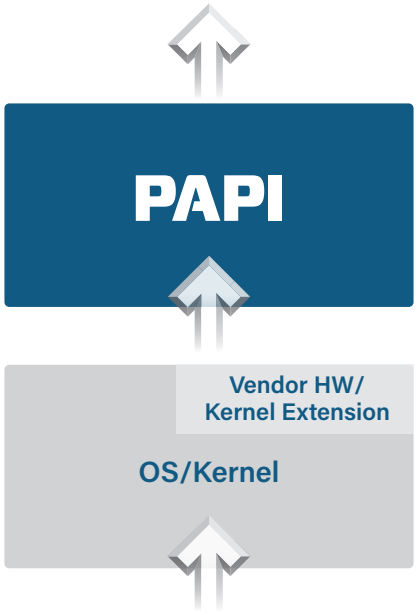


## SUPPORTED ARCHITECTURES

**AMD**
up to Zen4, power for Fam17h

GPUs MI50, MI60, MI100, M210, M250x, XGMI, power, temperature, fan

**ARM**
Cortex A7, A8, A9, A15, ARM64, ARM uncore-support

**Cray / HPE**
Interconnects: Gemini, Aris, Slingshot

RAPL power

**IBM**
Blue Gene Series, Q: 5D-Torus, I/O system, EMON power/energy

Power Series, PCP for POWER9-nest, power monitoring & capping

**Infiniband**

**Intel**
Nehalem, Westmere, Sandy Bridge, Ivy Bridge, Haswell, Haswell-EP, Broadwell, Sky-, Kaby-, Cascade-, Ice-lake, Sapphire Rapids, Knights Corner, Knights Landing, Knights Mill

RAPL; power capping, Power on Xeon Phi

Intel RAPL (power/energy), power capping capabilities (via PAPI_write())

GPUs: Ponte Vecchio

**Lustre**

**NVIDIA**
Tesla, Kepler, Maxwell, Pascal, Volta, Turing, Ampere, Hopper; support for multiple GPUs

support for NVLink

NVIDIA GPU support for power/energy reading; power capping capabilities (NVML)

NVIDIA PerfWorks API

**Virtual Environment**
VMWare, KVM


IN COLLABORATION WITH

THE UNIVERSITY OF MAINE

SPONSORED BY

U.S. Department of Defense
U.S. DEPARTMENT OF ENERGY
ECP
National Science Foundation

WITH SUPPORT FROM

AMD
CRAY THE SUPERCOMPUTER COMPANY
IBM
intel
NVIDIA

# EXA-PAPI


## Extending PAPI for ECP Applications

The Exa-PAPI project has developed new performance counter monitoring capabilities as well as power management support for novel and advanced ECP hardware, and software technologies. Exa-PAPI, building upon the functionality of classic-PAPI, has added performance and power monitoring support for AMD GPUs through integration with AMD ROCm and ROCm-SMI. It also supports Intel Ponte Vecchio GPUs via Intel's oneAPI Level Zero and NVIDIA GPUs through the Perfworks API. Additionally, Exa-PAPI accommodates the Slingshot interconnect, as well as the latest CPUs and ARM chips. All this new functionality was added while maintaining the same standard PAPI interface and methodology for using low-level performance counters in CPUs, GPUs, on/off-chip memory, interconnects, and the I/O system, including energy/power management. In addition to providing support for the latest hardware, Exa-PAPI improved integration with third-party tools, such as TAU, providing a portable and unified approach to performance and power monitoring through the PAPI library. ECP also improved PAPI's sustainability by enabling integration into Spack and E4S, and ensuring software robustness through continuous integration and continuous deployment (CI/CD). In addition to providing hardware counter-based information, the ability to register and monitor Software-Defined Events (SDE) has been added to PAPI. Through SDEs the internal behavior of runtime systems and libraries, such as PaRSEC, SLATE, Magma, Ginkgo, etc, are exposed to the applications that use those libraries.

As a result, Exa-PAPI has expanded the concept of performance events beyond strictly hardware-related events to include software-based information. Additionally, it has extended the range of supported hardware by including advanced performance and power monitoring capabilities for the latest CPUs, accelerators, and interconnects found in modern exascale platforms.

## Key Challenges

Widely deployed and widely used, PAPI has established itself as fundamental software infrastructure in every application domain where improving performance can be mission critical.

However, processor and system designs have been experiencing radical changes. Systems now combine multi-core CPUs and accelerators, shared and distributed memory, PCI-express and other interconnects, and power efficiency is emerging as a primary design constraint. These changes pose new challenges and bring new opportunities to PAPI. At the same time, the ever-increasing importance of communication and synchronization costs in parallel applications, as well as the emergence of task-based programming paradigms, pose challenges to the development of performance-critical applications and create a need for standardizing performance events that originate from various ECP software layers.

## Solution Strategy

The Exa-PAPI team extended PAPI to stand up to the challenges posed by exascale systems by:

- widening its applicability and providing robust support for exascale hardware resources;
- supporting finer-grain measurement and control of power, thus offering software developers a basic building block for dynamic application optimization under power constraints;
- extending PAPI to support software-defined events; and
- applying semantic analysis to hardware counters so that the application developer can better make sense of the ever-growing list of raw hardware performance events that can be measured during execution.

In summary, the team channeled the monitoring capabilities of hardware counters, power usage, software-defined events into a robust PAPI software package.

## RECENT PUBLICATIONS


Barry, D., Jagode, H., Danalis A., Dongarra, J.
**Memory Traffic and Complete Application Profiling with PAPI Multi-Component Measurements**
*2023 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, St. Petersburg, Florida, 2023

Barry, D., Danalis, A., Jagode, H.
**Effortless Monitoring of Arithmetic Intensity with PAPI's Counter Analysis Toolkit**
*13th International Workshop on Parallel Tools for High Performance Computing,* Springer Nature Switzerland AG, C. Niethammer et al. (Eds.), pp. 1-25, 2020.

Jagode, H., A. Danalis, H. Anzt, J. Dongarra
**PAPI Software-Defined Events for in-Depth Performance Analysis**
*The International Journal of High Performance Computing Applications* 33, no. 6 (November 2019): 1113–27. doi:10.1177/1094342019846287.

Danalis, A., H. Jagode, T. Herault, P. Luszczek, J. Dongarra.
**Software-defined Events through PAPI**
*24th International Workshop on High-Level Parallel Programming Models and Supportive Environments (HIPS), in conjunction with 33rd IEEE International Parallel & Distributed Processing Symposium (IPDPS),* May 20-24, 2019 , Rio de Janeiro, Brazil, pp. 1-10, 2019.



SPONSORED BY

U.S. DEPARTMENT OF ENERGY | Office of Science

This research was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration.


EXASCALE COMPUTING PROJECT
https://exascaleproject.org