# Grm

## Arm Forge for IO Profiling

Florent Lebeau 21/11/19

© 2019 Arm Limited

### Arm Forge Ultimate

A cross-platform toolkit for debugging, profiling and performance analysis



Commercially supported by Arm





The de-facto standard for HPC development

- Available on the vast majority of the Top500 machines in the world
- Fully supported by Arm on Arm servers, x86, IBM Power, Nvidia GPUs, etc.

State-of-the art debugging and profiling capabilities

- Powerful and in-depth error detection mechanisms (including memory debugging)
- Sampling-based profiler to identify and understand bottlenecks
- Available at any scale (from serial to petaflopic applications)

Easy to use by everyone

- Unique capabilities to simplify remote interactive sessions
- Innovative approach to present quintessential information to users

#### Arm Forge – MAP Multi-node Low-overhead Profiler

#### Inspect **OpenMP** activity



#### Usage

- Compile executable and libraries with debugging information \$ mpicc -g -O3 myapp.c -o myapp.exe
- Profile by prefixing your mpirun/srun command \$ map --profile mpirun -n 8 ./myapp.exe myargs \$ map --profile mpirun -n 8 python ./myapp.py myargs
- This will create a \*.map file that can be open in the GUI \$ map myapp\_8p\_1n\_YYYY-MM-DD\_HH-MM.map

#### Visualizing profiling results



### Visualizing profiling results



### Analyzing IO usage

- POSIX read/write rate, POSIX read/write syscall rate
  - Total IO rates from the application
- Disk read/write transfer rates
  - Includes disk and network filesystems accesses. May not include all IO due to page caching.

Preset: Default		
Preset: Activity Timelines	<u>F</u> ile <u>E</u> dit <u>Vi</u> ew <u>M</u> etrics <u>W</u> indow <u>H</u> elp	
Preset: CPU Instructions	Profiled: python2.7 on 8 processes, 1 node, 8 cores (1 per process) Sampled from: Thu Nov 14 2019 14:02:15 (UTC) for 37.5s	Hide Metrics
Preset: CPU Time	Main thread activity	
Preset: Energy		
Preset: IO	POSIX I/O read rate 345	
Preset: Linux perf CPU events	0.09 kB/s	
Preset: Memory	POSIX I/O write rate 14.8	
Preset: MPI	0.37 MB/s	and the second
A shirika Tina slip sa	Disk read transfer 27.0	
Activity Timelines	0.00 MB/s	
CPU Instructions	Disk write transfer	
CPU Time	0.37 MB/s	and the second se
Energy •		
10		
Linux perf CPU events		
Memory •	POSIX write syscall rate 161	
MPI 🕨	0.05 K calis/s	A little and an it will be an a first of the second
Unavailable groups	15:02:15-15:02:52 (37.451s): Main thread compute 43.3 %, MPI 53.5 %, File I/O 0.6 %, Python interpreter 2.6 %, Synchronisation 0.0 %, Sleeping 0.0 %	zoom 🔍 🇮 🕥

#### Lustre metrics

- Query the kernel for Lustre data activity
- Stores read / write rates, volume, file opens and metadata activity



#### Arm Performance Reports Application Analysis Tool



#### arm

## **Extending IO Metrics - NVRAM**

- Comprehensive set of NVRAM metrics
  - Supporting 1LM and 2LM mode
- Memory bandwidth data
  - Read and write
  - For each socket
- Load instructions
  - NVRAM and DDR
- Memory consumption
  Allocation on NVDIMMs
- Compare against existing metrics
  - E.g. DRAM memory consumption

Profiled: pstreams on 1 process, 1 node, 8 cores (8 per process) Sampled from: Fri Sep 27 2019 12:29:



L 1	L J	L .	L	L .	L .	L	L .	L .	L .	±	L .	L .	L J	L.
	F 7													

	rn	$\mathbf{n}^{+}$					Thank You
+	+	+					Danke
							Merci
							谢谢
							ありがとう
							Gracias
							Kiitos
							감사합니다
							धन्यवाद
							شکرًا
							תודה
© 2019 Arm	Limited						

+ +	+ + +	+ +	+ + +	+ -	- + +	

			The Arm trademarks featured in this presentation are registered trademarks or trademarks of Arm Limited (or its subsidiaries) in the US and/or elsewhere. All rights reserved. All other marks featured may be trademarks of their respective owners
			reatured may be trademarks of their respective owners.

www.arm.com/company/policies/trademarks

© 2019 Arm Limited