



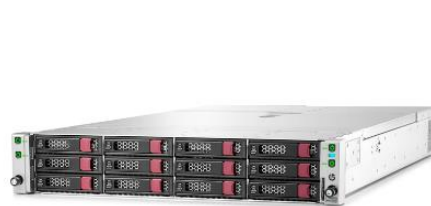
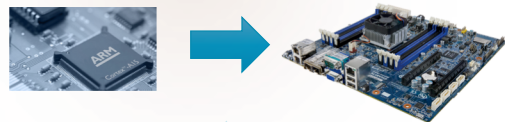
arm

Open UCX on Arm

Pavel Shamis / Pasha
Principal Research Engineer

SC 2018

Scaling



• **HPE Apollo 70:** This dense, scalable platform – and HPE's first ARM-based HPC system – brings more choice and flexibility to HPC customers. It provides easy access to HPC technology with support for standard HPE provisioning, cluster management and performance software. The Apollo 70, using Cavium's 64-bit ARMv8-A ThunderX2™ Server Processor, is purpose-built for memory intensive HPC workloads and delivers up to 33 percent more memory bandwidth than today's industry standard servers. The Apollo 70 also provides access to HPE's partnership ecosystem delivering key HPC components including Red Hat® Enterprise Linux®, SUSE® Linux Enterprise Server for ARM, and Mellanox® high-speed InfiniBand & Ethernet fabric solutions.

University of
BRISTOL

GW⁴

January 17th 2017

Announcing the **GW4 Tier 2 HPC service, 'Isambard'**:
named after Isambard Kingdom Brunel

System specs:

- Cray CS-400 system
- **10,000+ ARMv8 cores**
- HPC optimised software stack
- Technology comparison:
 - x86, KNL, Pascal
- To be installed March-Dec 2017
- £4.7m total project cost over 3 years

Simon McIntosh Smith, simonm@cs.bris.ac.uk, @simonm

Sandia's NNSA/ASC ARM Platforms

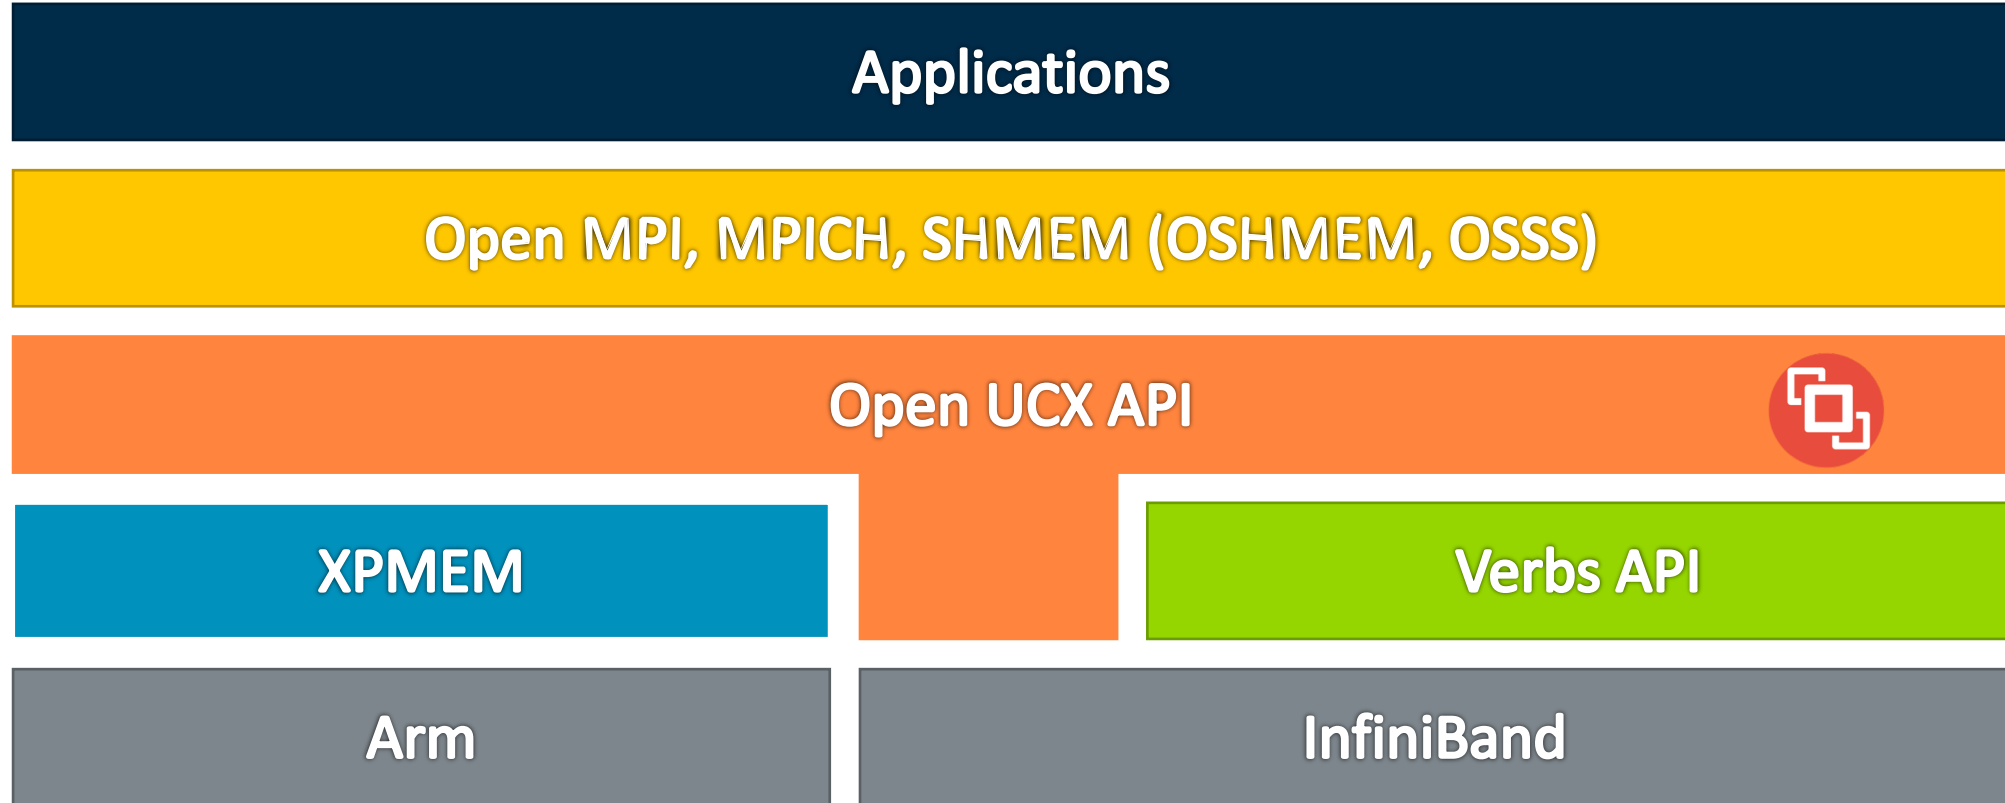


Open UCX on Arm



- Works seamlessly on Arm v8
 - <https://developer.arm.com/products/software-development-tools/hpc/resources/porting-and-tuning/building-openmpi-with-openucx>
- uArch level optimizations for ThunderX2 and other Arm implementations
- Optimized for NEON SIMD
- Optimized memory pulling routines for Arm
- Tested with multiple Arm v8 implementations
- Tested with InfiniBand, ROCE, TCP, shared memory

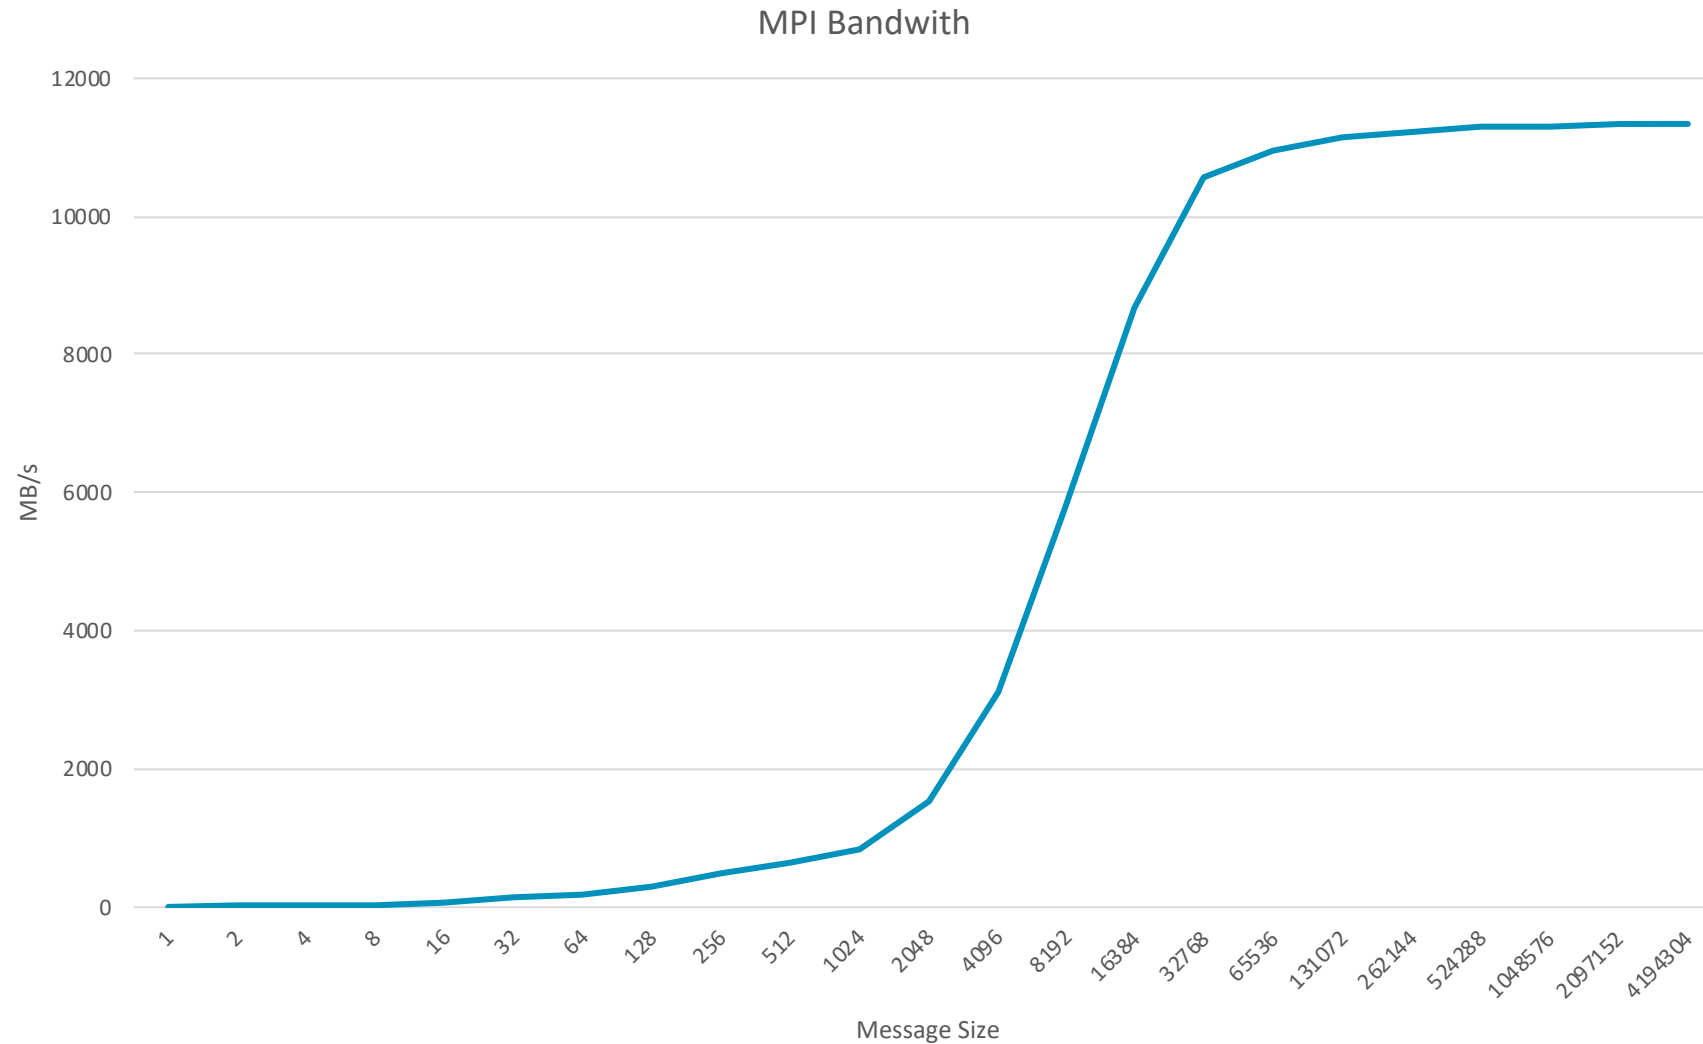
Hardware Software Stack Example



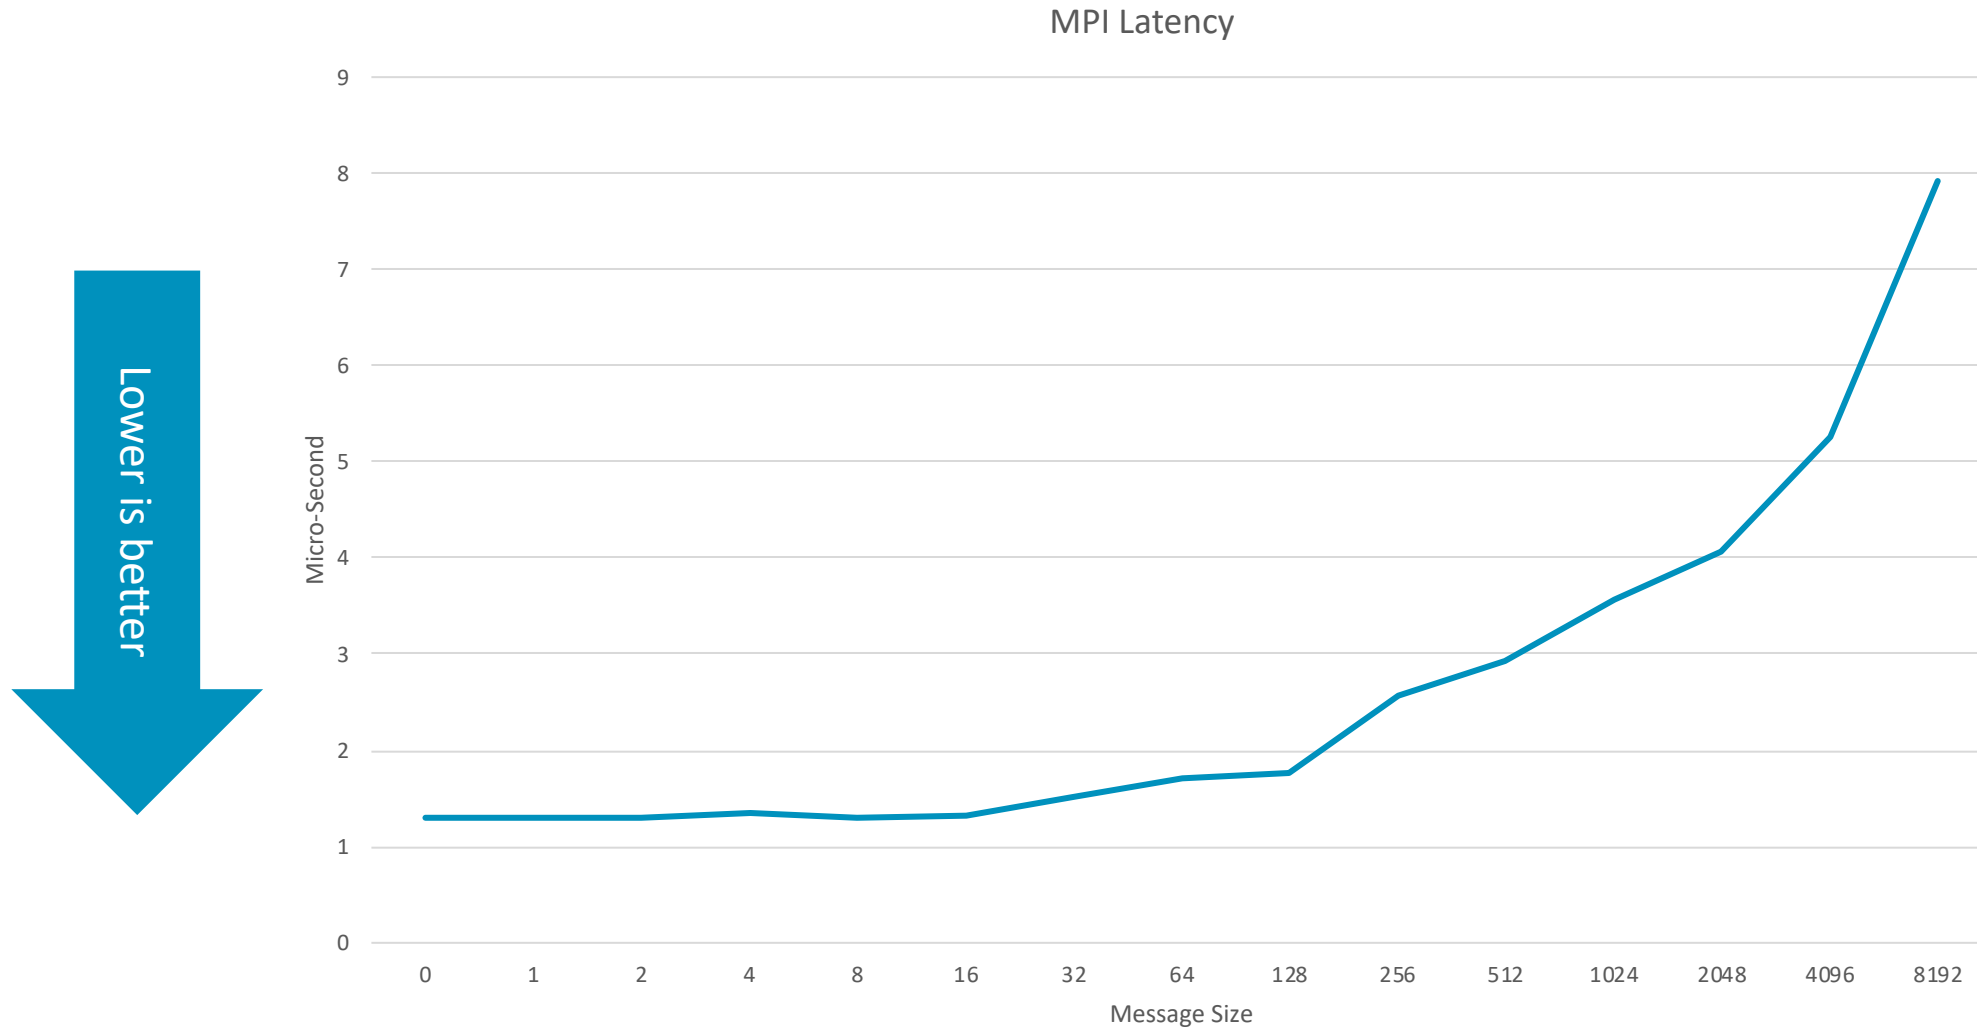
Testbed

- HPE Comanche platform with Cavium ThunderX2 (pre-production)
- Mellanox ConnectX-4 EDR (100Gb/s) HCA
- Mellanox 100Gb/s Switch
- OpenMPI 3.0 with OpenUCX 1.3 (MLNX_OFED_LINUX-4.3-1.0.1.0)
- OSU Micro-Benchmarks

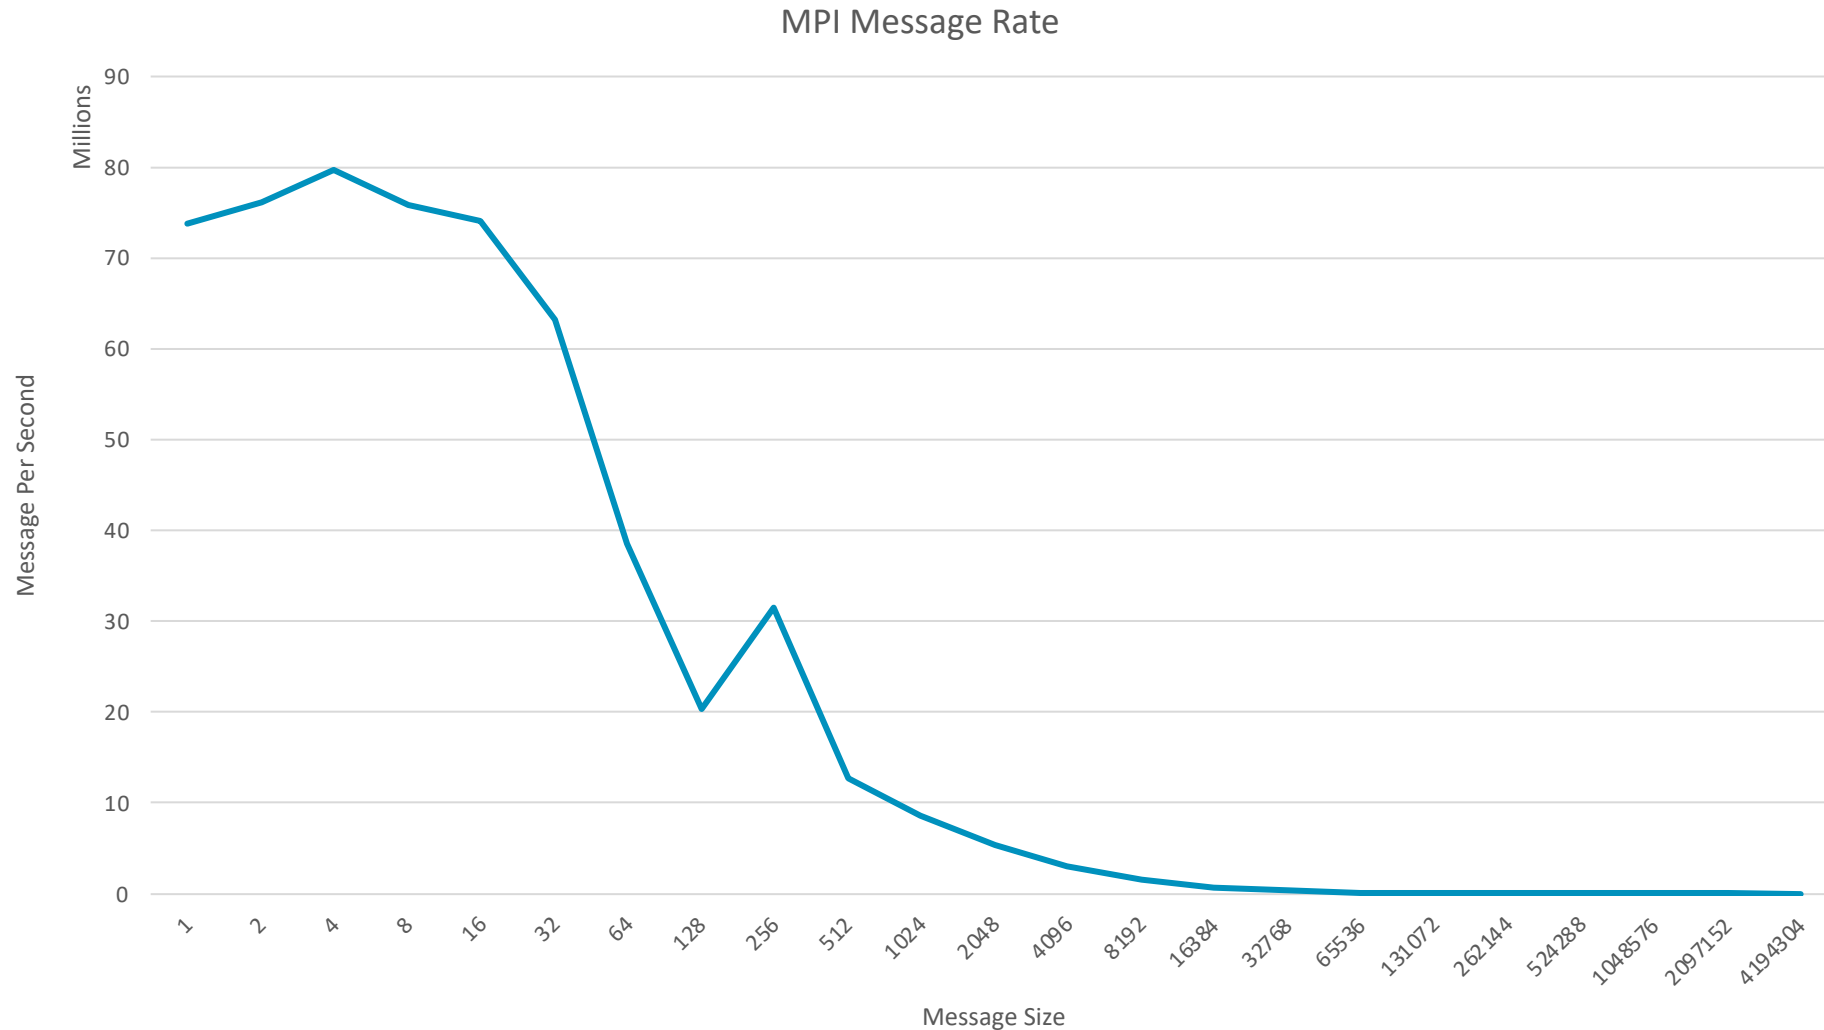
Cavium Thunder X2 **SINGLE** core InfiniBand Bandwidth



Cavium Thunder X2 MPI Ping-Pong Latency with InfiniBand



Cavium Thunder X2 MPI Message Rate with InfiniBand (28 cores)



Higher is better

The word "arm" in a white, lowercase, sans-serif font, positioned on the right side of a solid blue background.

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.

www.arm.com/company/policies/trademarks

SHMEM_WAIT()

