

E4

**COMPUTER
ENGINEERING**
INT 21H
POP AX

I've got the Power!

*Performance comparison between state of the art x86_64 and
OpenPOWER servers for High Energy Physics and related
considerations*

Daniele Gregori

daniele.gregori@e4company.com

HPC & ENTERPRISE SOLUTIONS

E4® Computer Engineering SpA was founded in 2002



We are specialized in the **manufacturing of high performance IT systems** of medium and high range. Our products aim to meet both **industrial and scientific research requirements** with range fit for many environments, from universities to computing centers.

Thanks to the established experience acquired through the years, E4 is a valued technology's supplier acknowledged and appreciated by **many worldwide organizations.**

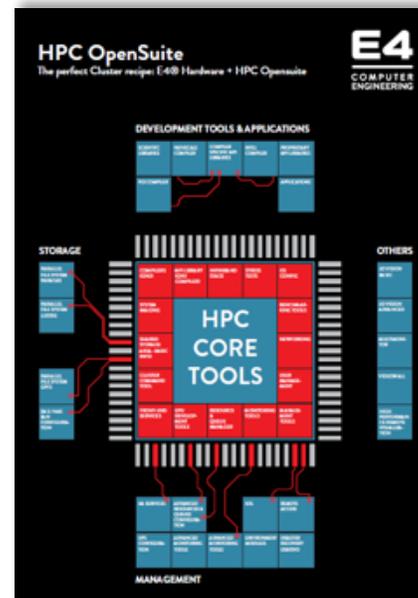
HPC Team

- Operating since 2005
- E4 HPC custom suite
- Unique technical skills in Italy
- E4 on SPEC website
- Intel Cluster Ready certified
- SGI fully trained and qualified



Enterprise Team

- Born in 2010 - Operating since 2011
- New certified servers and storages
- Focus on «Business Continuity»
- E4 VSTONE (first italian complete virtualization solution)



E4 VSTONE
STRUCTURED VIRTUALIZATION SOLUTION



More than 15M euros of installed HW

Requirements

High density computational nodes
Big data storage

Application

High-energy physics research

Hardware installed

- > 20PB high performance storage
- > 5PB direct attached storage
- > 3.500 server dual socket (~ 40k computing cores)

Several GPU systems
NBD intervention times

INFN
Istituto Nazionale
di Fisica Nucleare

Riconoscimento di eccellenza nella collaborazione industriale per
gli esperimenti ATLAS e CMS al Large Hadron Collider del CERN in
occasione della scoperta del Bosone di Higgs.

E4 Computer Engineering S.p.A.
Scandiano (RE)

Fornitura di processori per il trigger di alto livello
e per i sistemi di calcolo T1/T2 degli esperimenti
ATLAS e CMS.

The CMS logo consists of the letters 'CMS' in a blue, sans-serif font, positioned above a stylized graphic of a particle detector's cross-section with colorful lines representing particle tracks.The ATLAS EXPERIMENT logo features the text 'ATLAS EXPERIMENT' in a blue, sans-serif font, with 'ATLAS' on the top line and 'EXPERIMENT' below it. To the left of the text is a stylized graphic of a particle detector's cross-section with colorful lines representing particle tracks.



More than 20M euros of installed HW

Requirements

- High density computational nodes
- Big data storage

Application

High-energy physics research

Solution

- 8.000+ dual socket mainboards (54.000+ cores)
- 50.000+ enterprise class hard disks (120PB Storage)

An advertisement graphic for E4 Computer Engineering. The background is black. At the top, it says "E4 SERVING UP HIGH ENERGY PHYSICS" in white. Below that, it lists "Take 7,200 SERVERS with 108,000 tasty CORES" and "Add 89 PB of STORAGE made of 42,900 crispy HDDs". At the bottom, it says "Season with ITALIAN PASSION" and "... AND THE KITCHEN IS STILL OPEN...". On the right, there is an illustration of a frying pan with a wooden handle, and several hard drive icons are shown as if they are being cooked. The E4 logo and "COMPUTER ENGINEERING" are in the top right, and "COOKING UP FAST COMPUTING" and "YEARS/ANS CERN" are in the middle right.

E4 SERVING UP HIGH ENERGY PHYSICS

Take **7,200 SERVERS**
with **108,000 tasty CORES**

Add **89 PB of STORAGE**
made of **42,900 crispy HDDs**

Season with **ITALIAN PASSION**

... AND THE KITCHEN IS STILL OPEN...

E4
COMPUTER
ENGINEERING

COOKING UP FAST COMPUTING

YEARS/ANS CERN

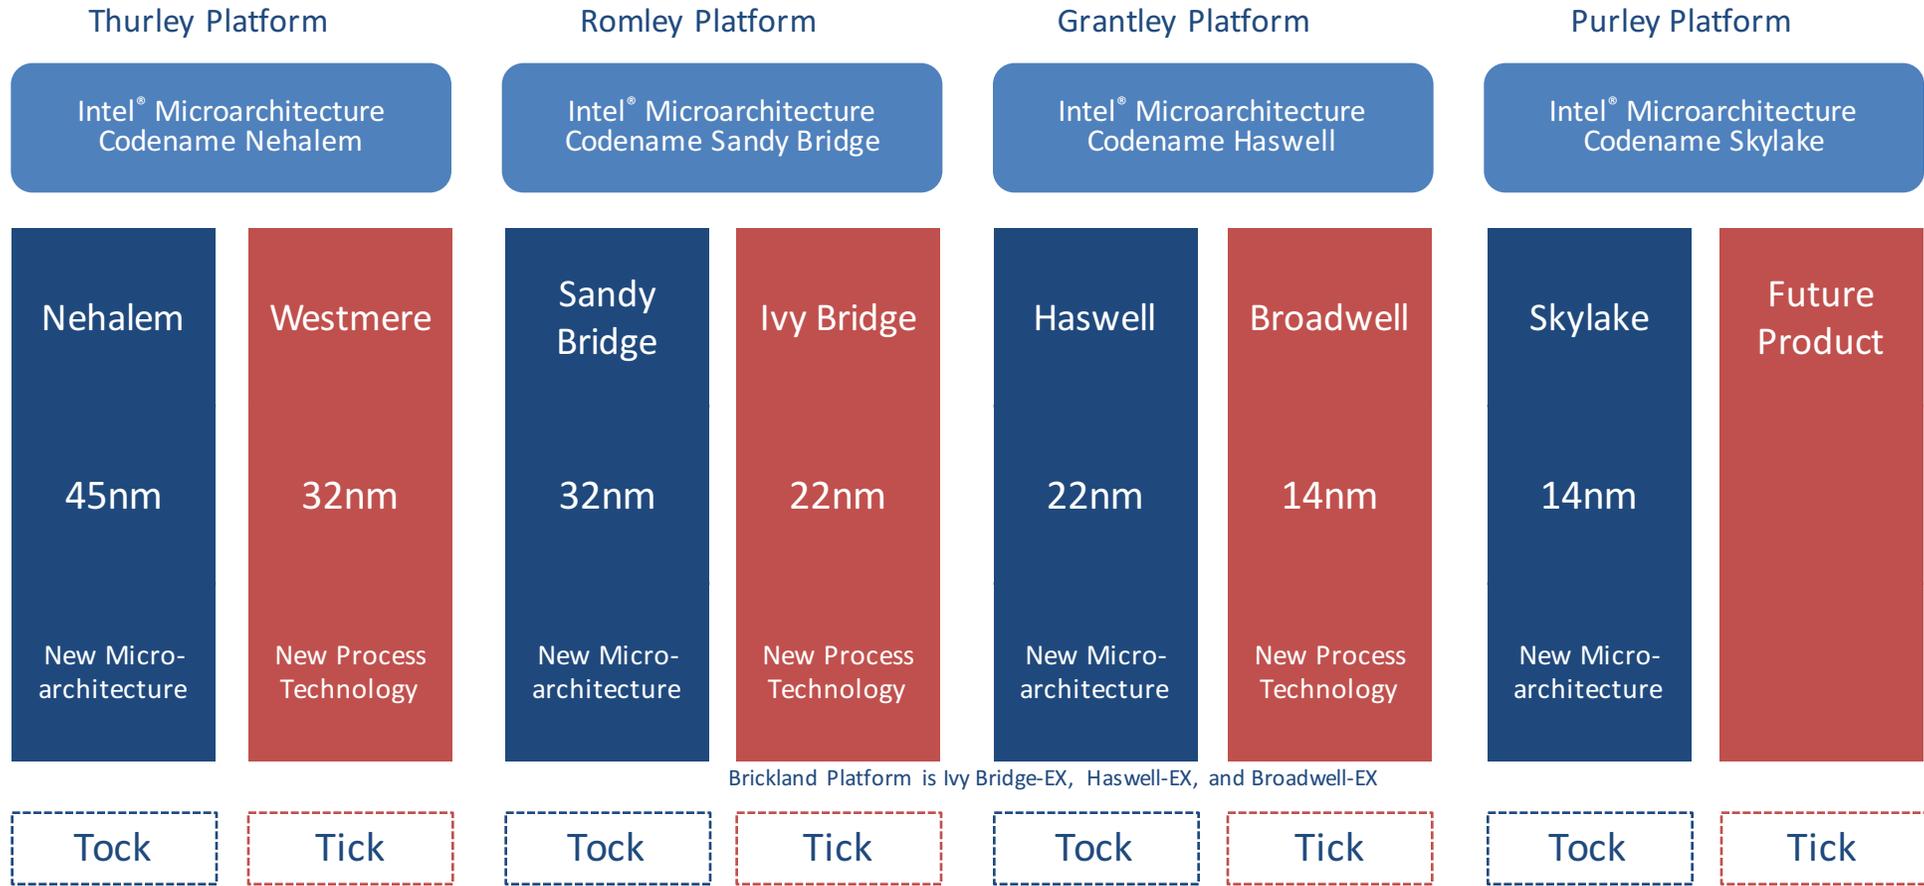
updated to 2014



E4® systems @ CERN

Tick-Tock Development Model:

Sustained Microprocessor Leadership HPC & ENTERPRISE SOLUTIONS

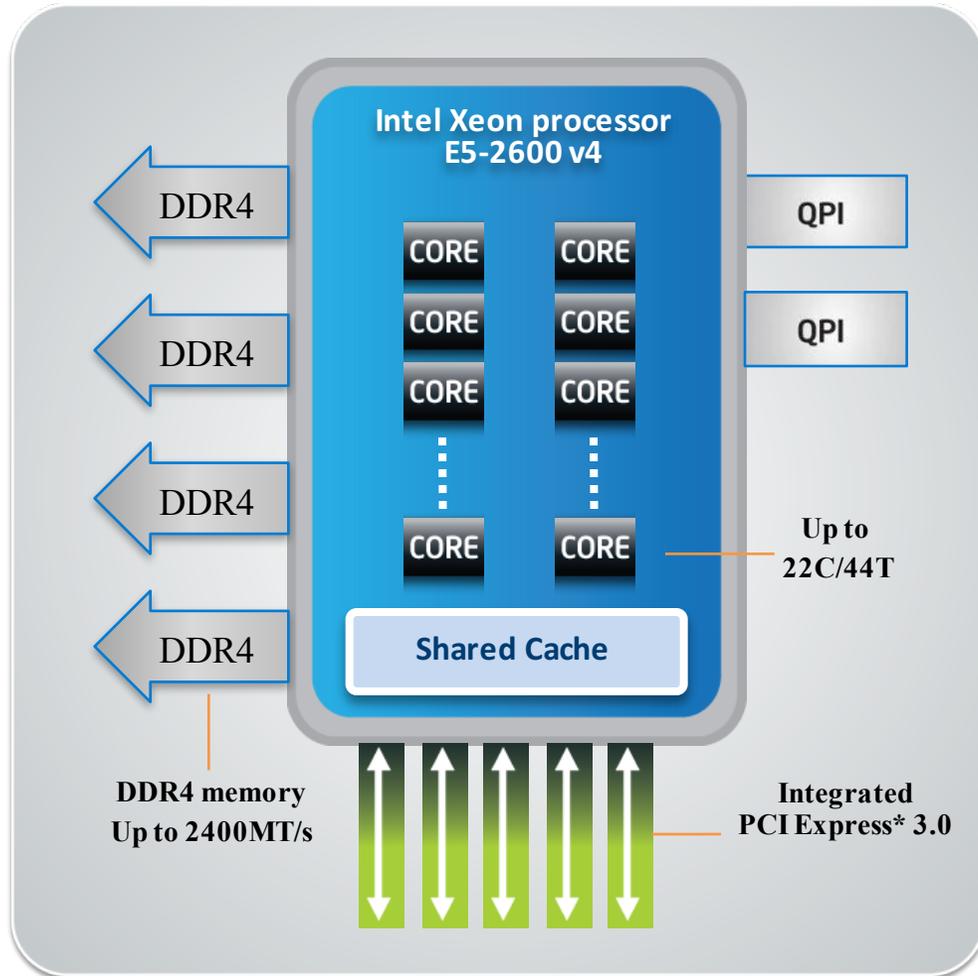


Brickland Platform is Ivy Bridge-EX, Haswell-EX, and Broadwell-EX

Innovation delivers new microarchitecture with Skylake

INTEL® XEON® PROCESSOR E5-2600 V4 PRODUCT FAMILY

HPC & ENTERPRISE SOLUTIONS



HIGHLIGHTS:

- Broadwell-EP processor based on 14nm process technology.
- Haswell-EP microarchitecture with new instructions including AVX2 and FMA.
- Up to 22C/44T DDR4 support increased memory speed at 2400MT/s.
- Compatible with Intel® C610 series chipset.

PLATFORM STORAGE EXTENSIONS:

- Asynchronous DRAM Refresh (ADR)
- PCIe Non-Transparent Bridge (NTB)
- Intel® QuickData Technology (CBDMA) & RAID-5 acceleration engine
- PCIe Dual-Cast

All products, computer systems, dates and figures specified are preliminary based on current expectations, and are subject to change without notice.

Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. Click http://www.intel.com/products/processor_number for details

OpenPOWER – OP205



OP205
A MOST POWERFUL CHOICE



HIGHLIGHTS

- High performance Linux server
- Dual IBM POWER8 processor modules
- Dual NVIDIA GPU accelerators
- Incorporates Mellanox Scalable HPC
- Solutions with NVIDIA technology
- CAPI technology
- Supports up to 1 TB of 1333/1066 MHz DDR3L memory
- Flexible and modular I/O
- Up to 8 threads
- Max 230GBps per socket

| FEATURES | OP205 |
|--------------------|--|
| Form Factor | 2U |
| Dimensions | D x W x H (mm) 730 x 442.5 x 86 |
| CPU | 2-socket, 8-core or 10-core IBM POWER8 processor/ Up to 190 W TDP per processor |
| GPU | Up to 2x NVIDIA K80 GPU accelerators |
| Memory | Up to eight memory riser modules; Up to 32 x RDIMM slots; 4 GB, 8 GB, 16 GB, and 32 GB 1333/1066 MHz DDR3 module ; 32 GB to 1024 GB capacity |
| I/O | 1 x USB 2.0 port (front); 1 x USB 3.0 port (rear); RJ45 port (dedicated for IPMI); 9-pin serial port; 15-pin VGA port |
| Power Supply | 2 x 1300W 80 PLUS Platinum hot-swap redundant power supplies (non-redundant if GPU card is installed) |
| Storage | 2 x storage bays for 2.5-inch SATA-III hard disk drive, hot-swappable, RAID 0, 1, 10 software implementation |
| Expansion Slot | One x16 PCIe Gen3 LP slot (CAPI enabled); Two x16 PCIe Gen3 FHFL riser cards (CAPI enabled); Two x8 PCIe Gen3 LP slot |
| System Management | IPMI 2.0 compliant ; AST 2400 CHIPSET |
| Storage Controller | SATA-III 6.0Gb/s ports (Marvell 88SE9235) |

POWER8 Processor

Technology

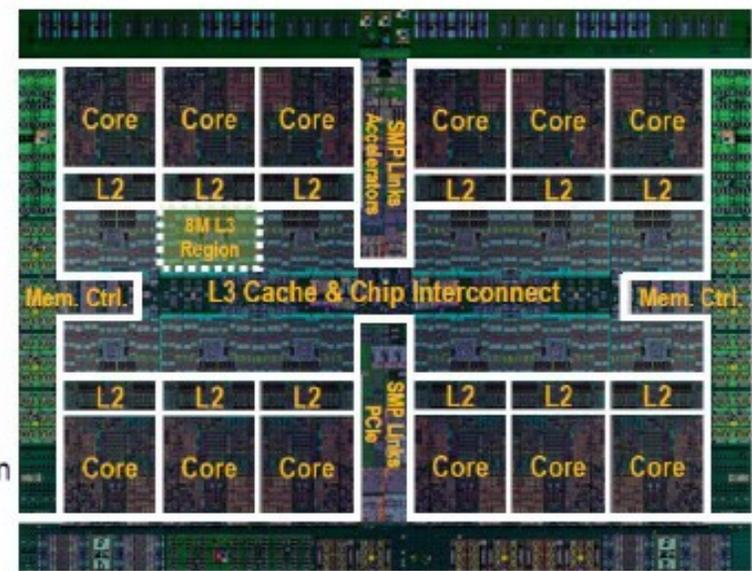
- 22nm SOI, eDRAM, 15 ML 650mm²

Cores

- 12 cores (SMT8)

Accelerators

- Crypto & memory expansion
- Transactional Memory
- VMM assist
- Data Move / VM Mobility



Energy Management

- On-chip Power Management Micro-controller
- Integrated Per-core VRM
- Critical Path Monitors

Caches

- 512 KB SRAM L2 / core
- 96 MB eDRAM shared L3
- Up to 128 MB eDRAM L4 (off-chip)

Memory

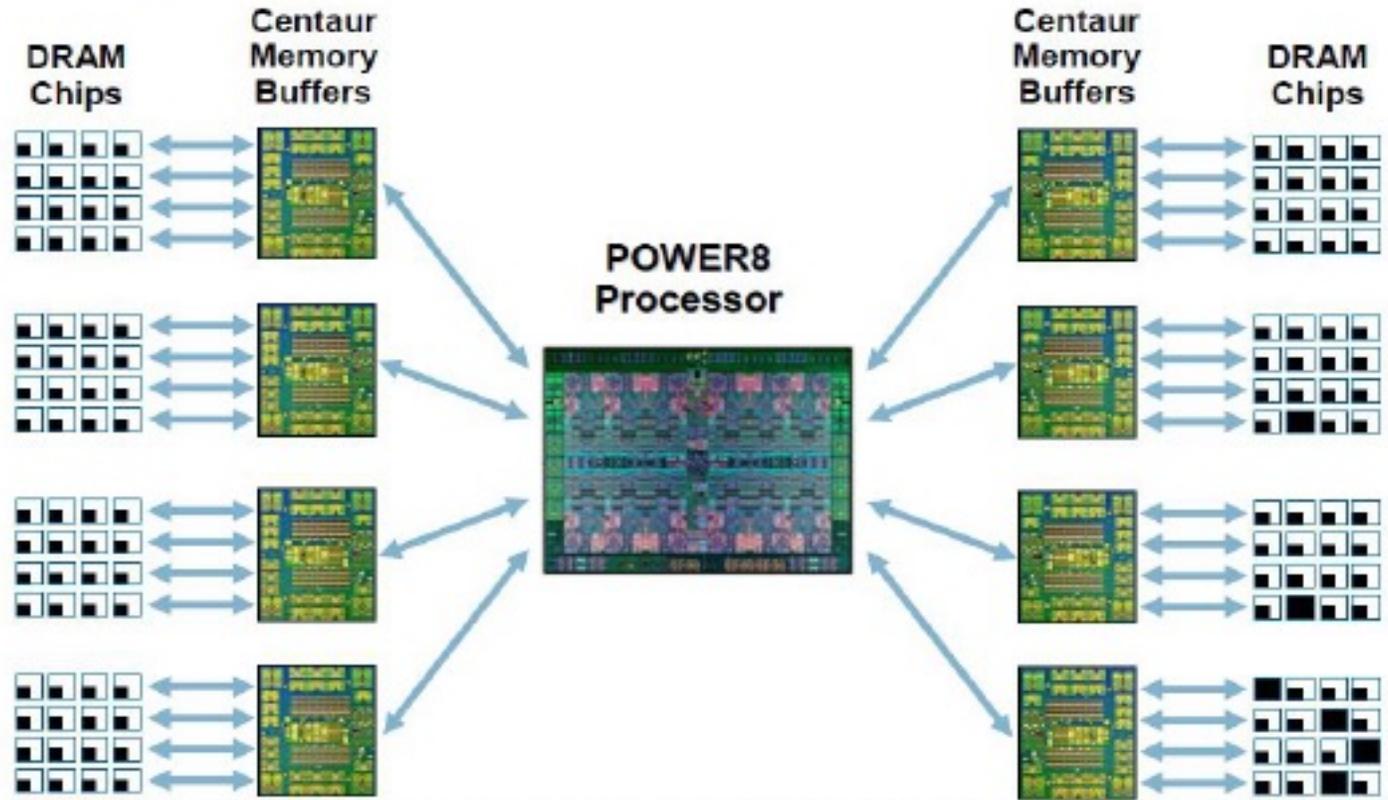
- Up to 230 GB/s sustained bandwidth

Bus Interfaces

- Durable open memory attach interface
- Integrated PCIe Gen3
- SMP Interconnect
- CAPI (Coherent Accelerator Processor Interface)



Memory Organization



- **Up to 8 high speed channels, each running up to 9.6 Gb/s for up to 230 GB/s sustained**
- **Up to 32 total DDR ports yielding 410 GB/s peak at the DRAM**
- **Up to 1 TB memory capacity per fully configured processor socket**



OpenPower8 2CPU 8core per CPU @ 3.8GHz, 512 GB RAM

OS: CentOS 7.2.15.11 ppc64le Compiler: gcc 4.8.5

Intel(R) Xeon(R) CPU E5-2697A v4 @ 2.60GHz, 128 GB RAM

OS: CentOS 6.6 Compiler: gcc 4.4.7

- Test HEPspec @ 64 bit
- Needs to recompile hepspec and toolset

HEPSPEC is the test adopted in HEP community to addresses the common workload

<https://w3.hepik.org/benchmarks/doku.php>

<https://www.spec.org/>

HEP SPEC

444.namd
447.dealll
450.soplex
453.povray
471.omnetpp
473.astar
483.xalancbmk

```
daniele.gregori — root@opower03:/opt/E4specCPU2006v1.2 — ssh root@100.100.3.5 — 117x
Installing FROM /opt/E4specCPU2006v1.2
Installing TO /opt/E4specCPU2006v1.2

Is this correct? (Please enter 'yes' or 'no')
yes

The following toolset is expected to work on your platform. If the
automatically installed one does not work, please re-run install.sh and
exclude that toolset using the '-e' switch.

The toolset selected will not affect your benchmark scores.

ppc64le-linux          Come se fosse Antani arch

=====
Attempting to install the ppc64le-linux toolset...

Checking the integrity of your source tree...
```

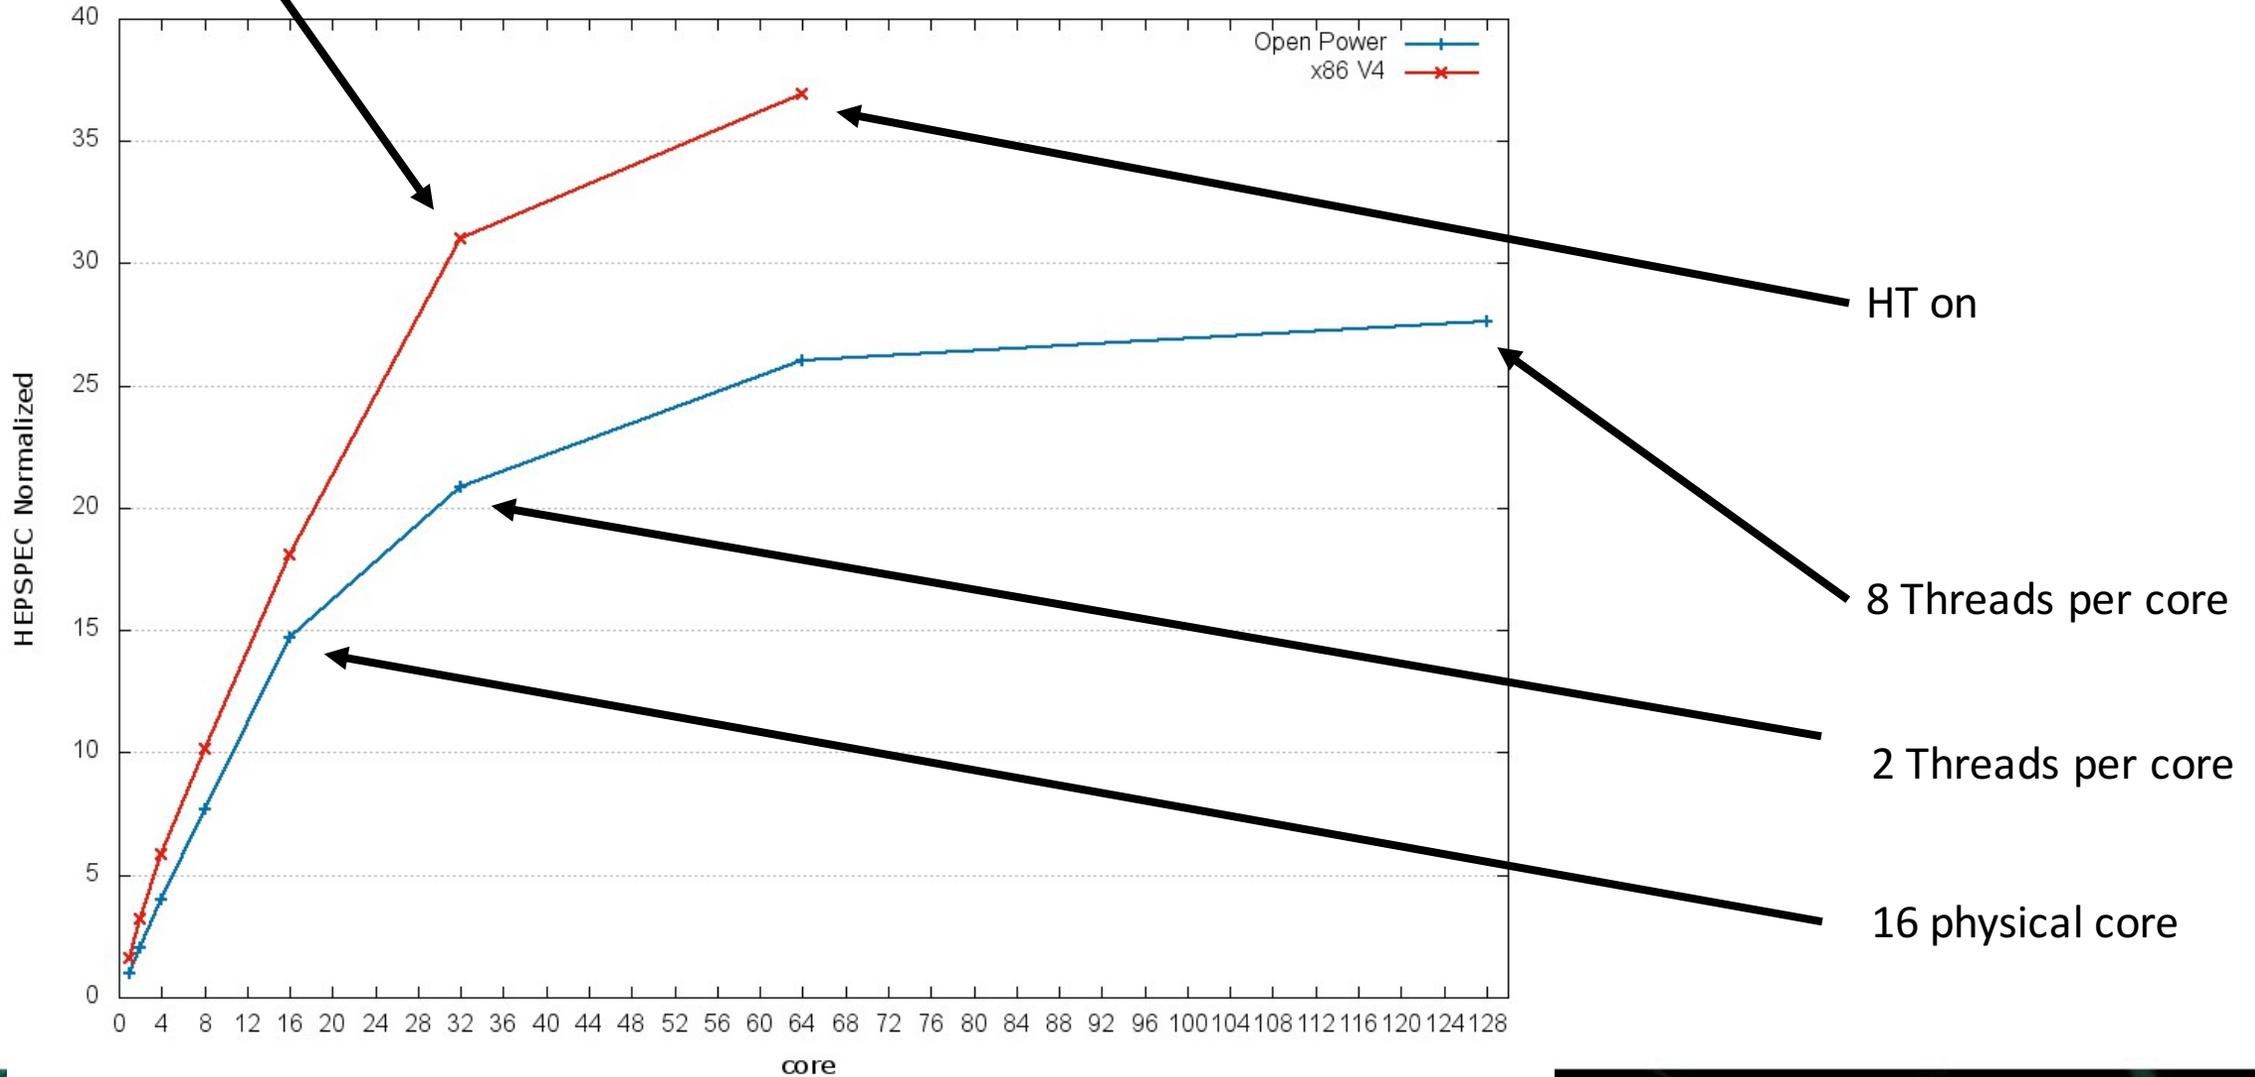
Results

HPC & ENTERPRISE SOLUTIONS



32 physical core

HEPSPEC per core: Open Power vs x86 V4 Normalized

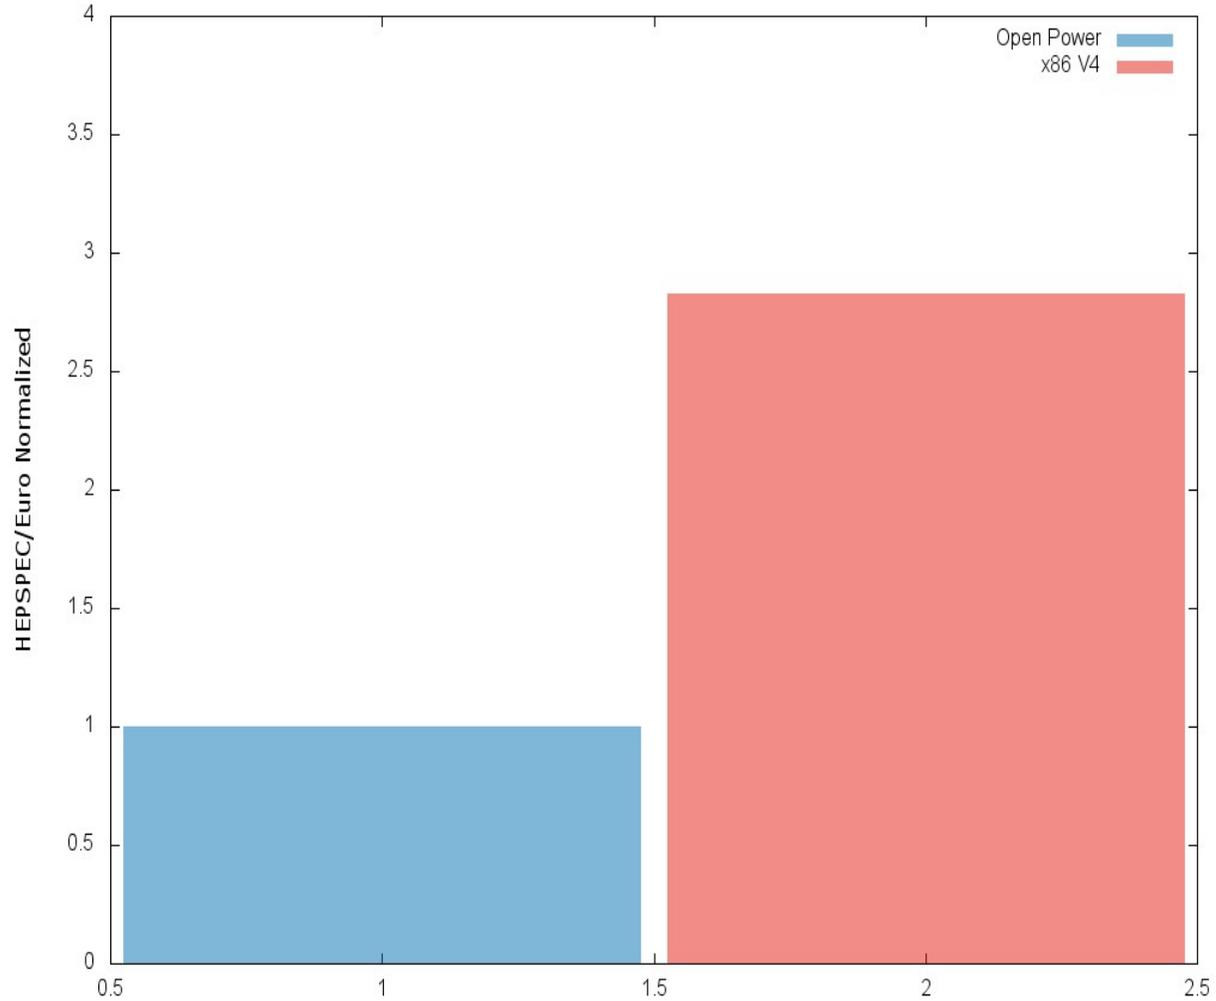


Results HEPSEC/Euro

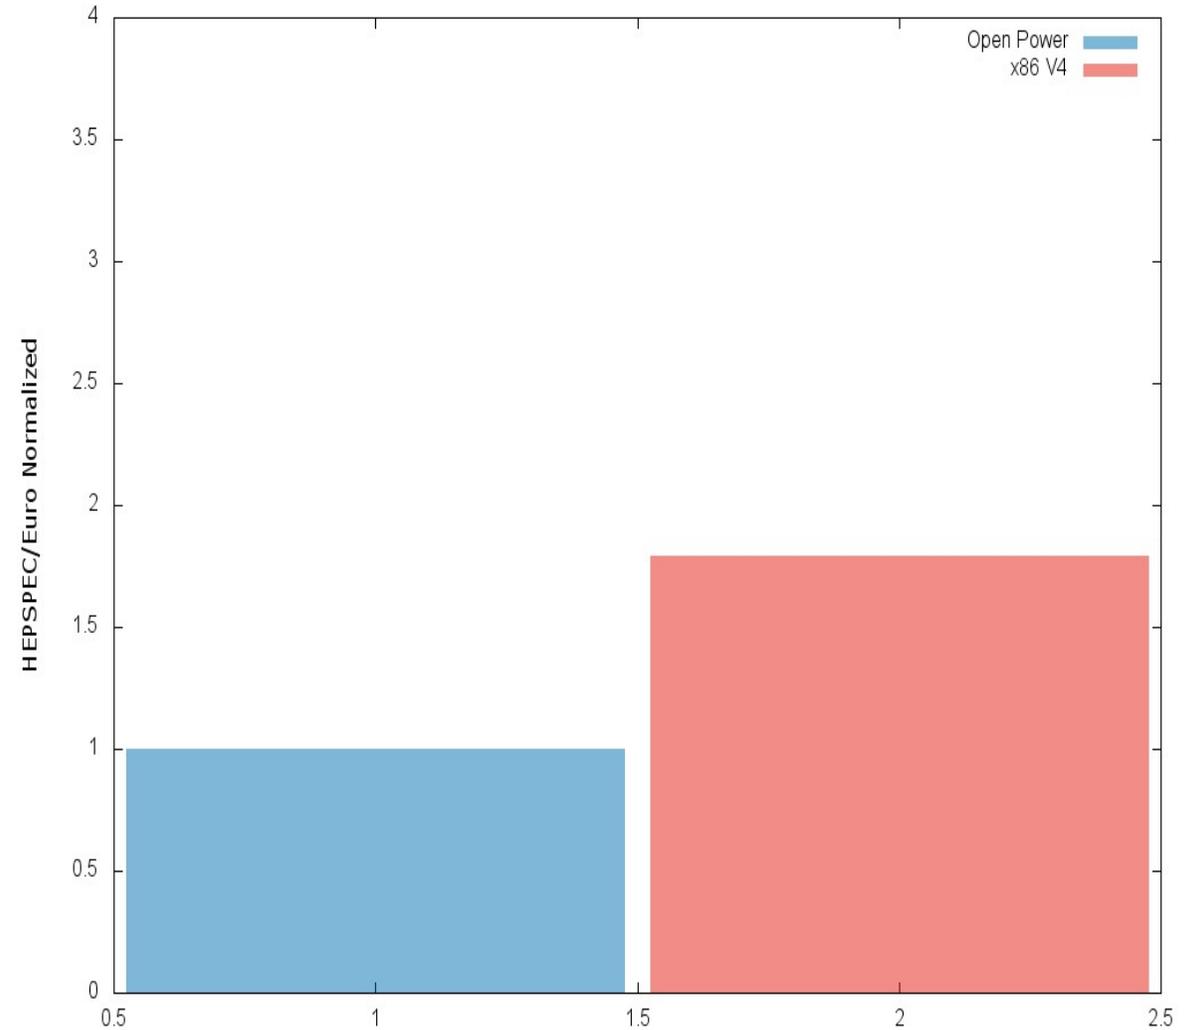
HPC & ENTERPRISE SOLUTIONS



HEPSPEC/Euro per CPU Normalized: Open Power vs x86 V4



HEPSPEC/Euro per max Threads Normalized: Open Power vs x86 V4

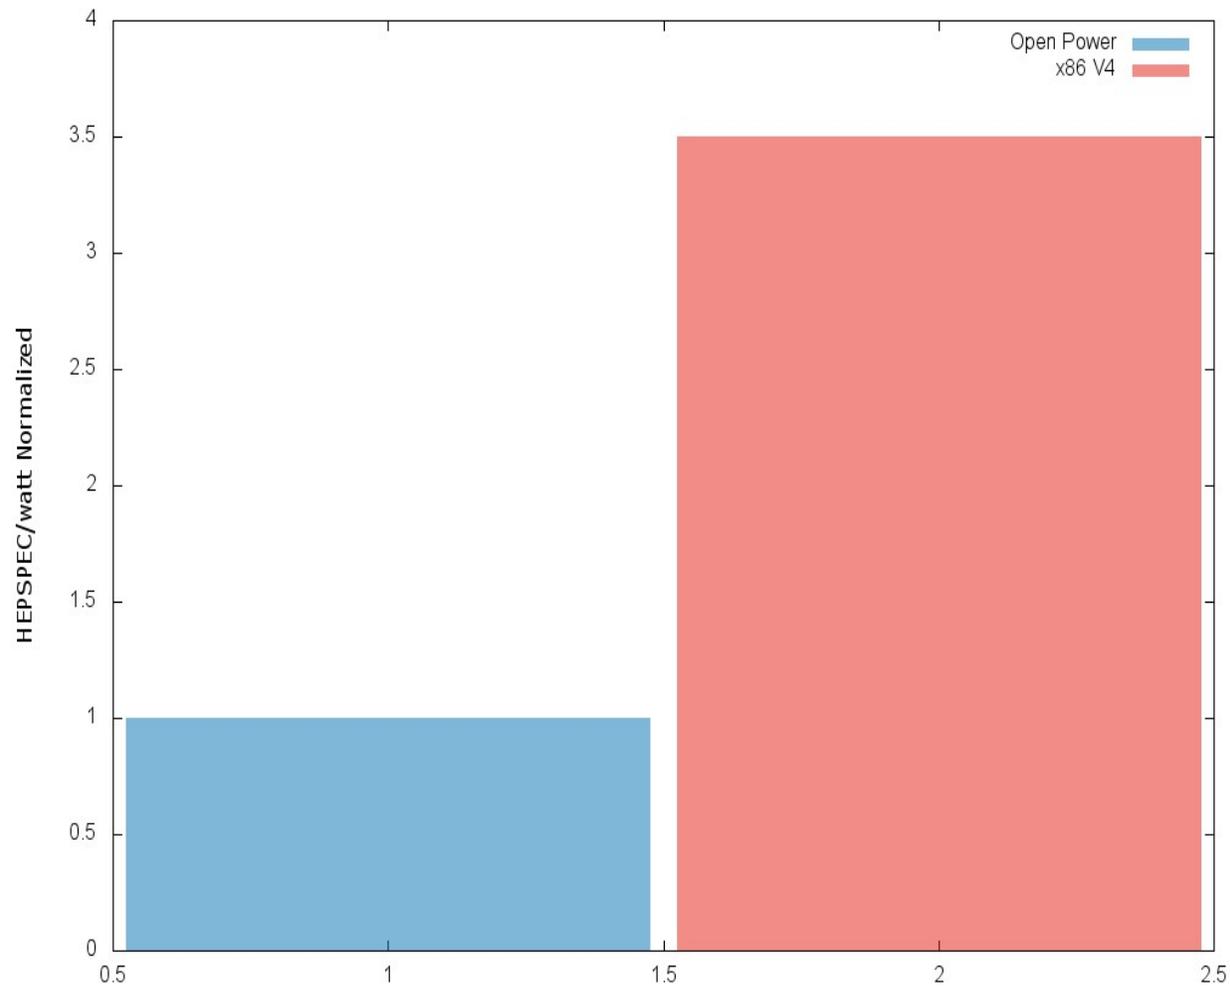


Results HEPSPec/watt

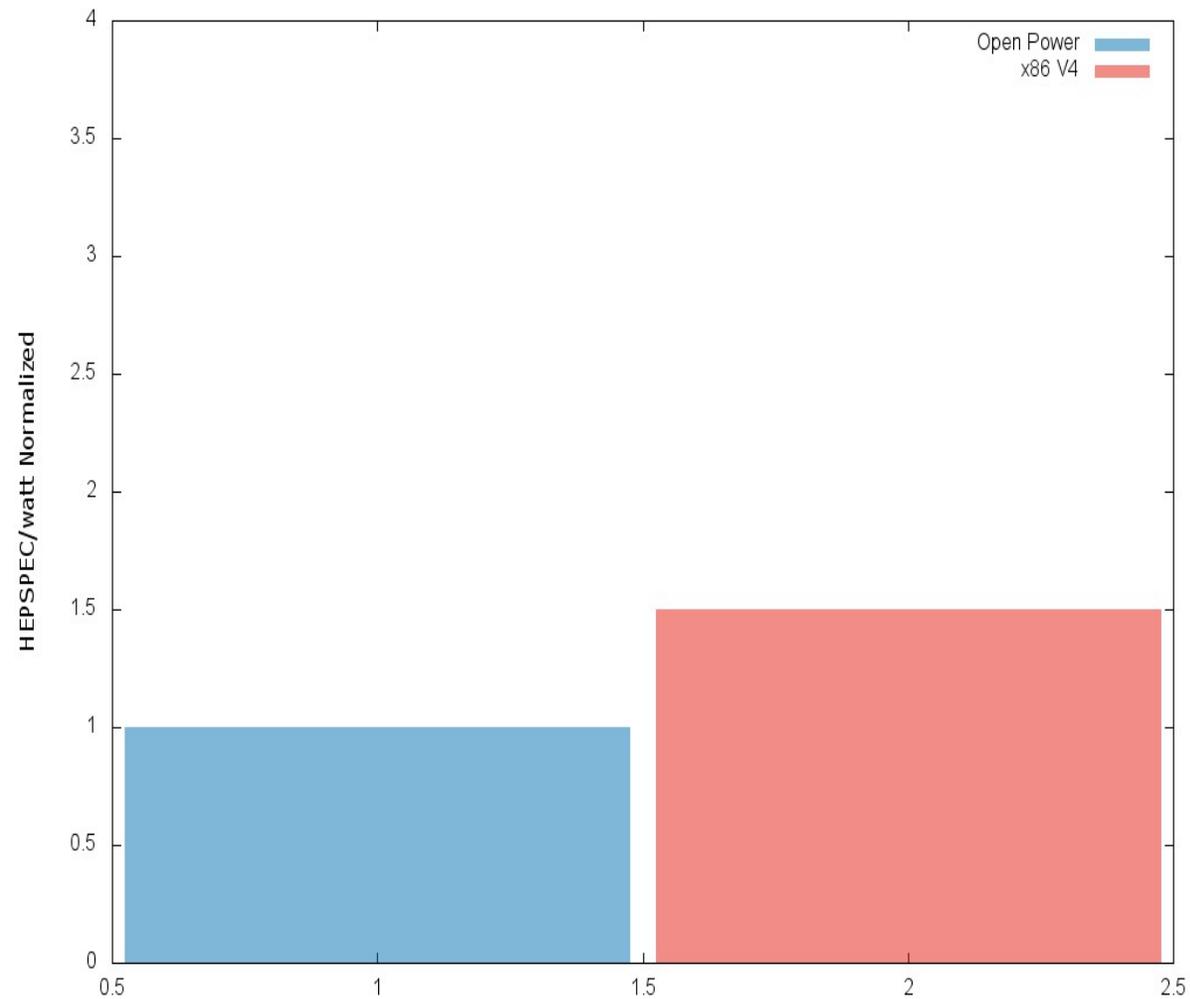
HPC & ENTERPRISE SOLUTIONS



HEPSPEC/watt per CPU Normalized: Open Power vs x86 V4



HEPSPEC/watt per max Threads Normalized: Open Power vs x86 V4

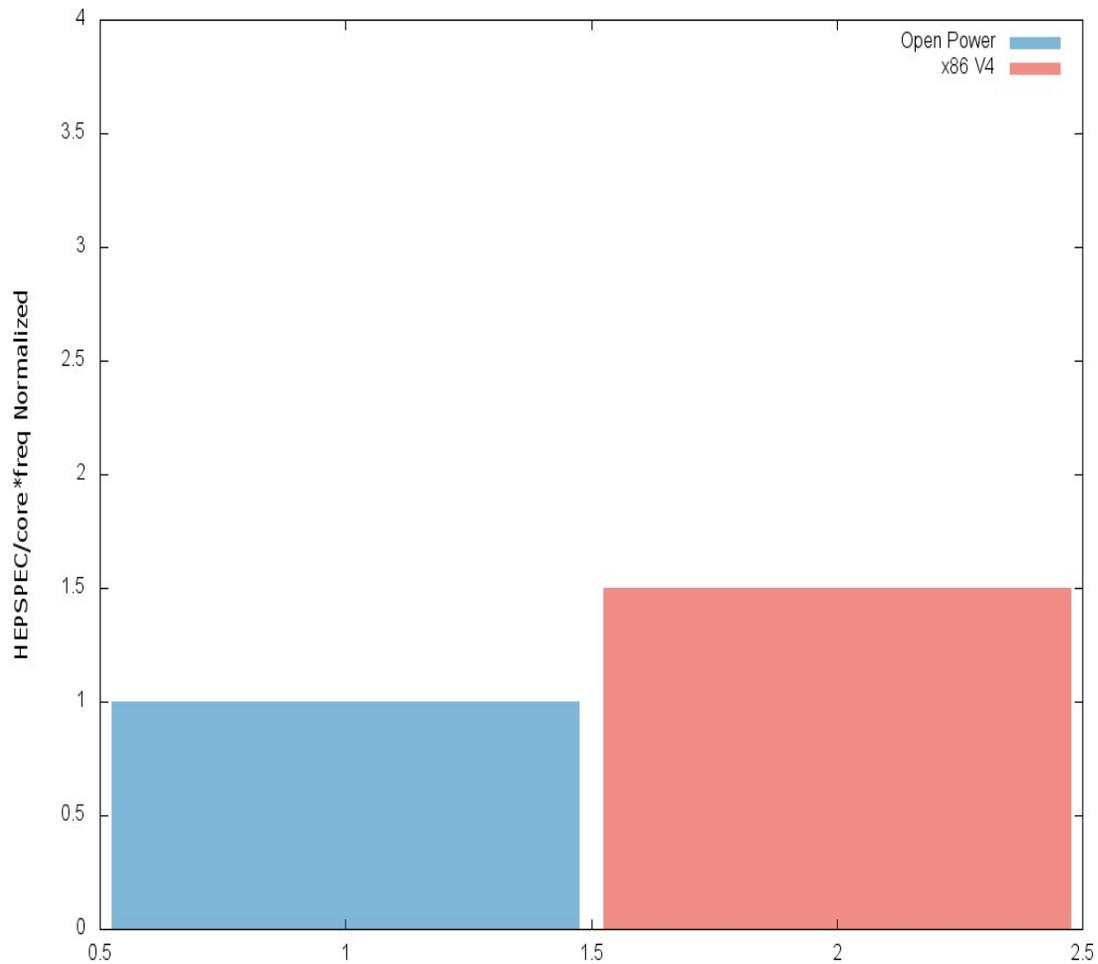


Results HEPSPREC/(threds*freq)

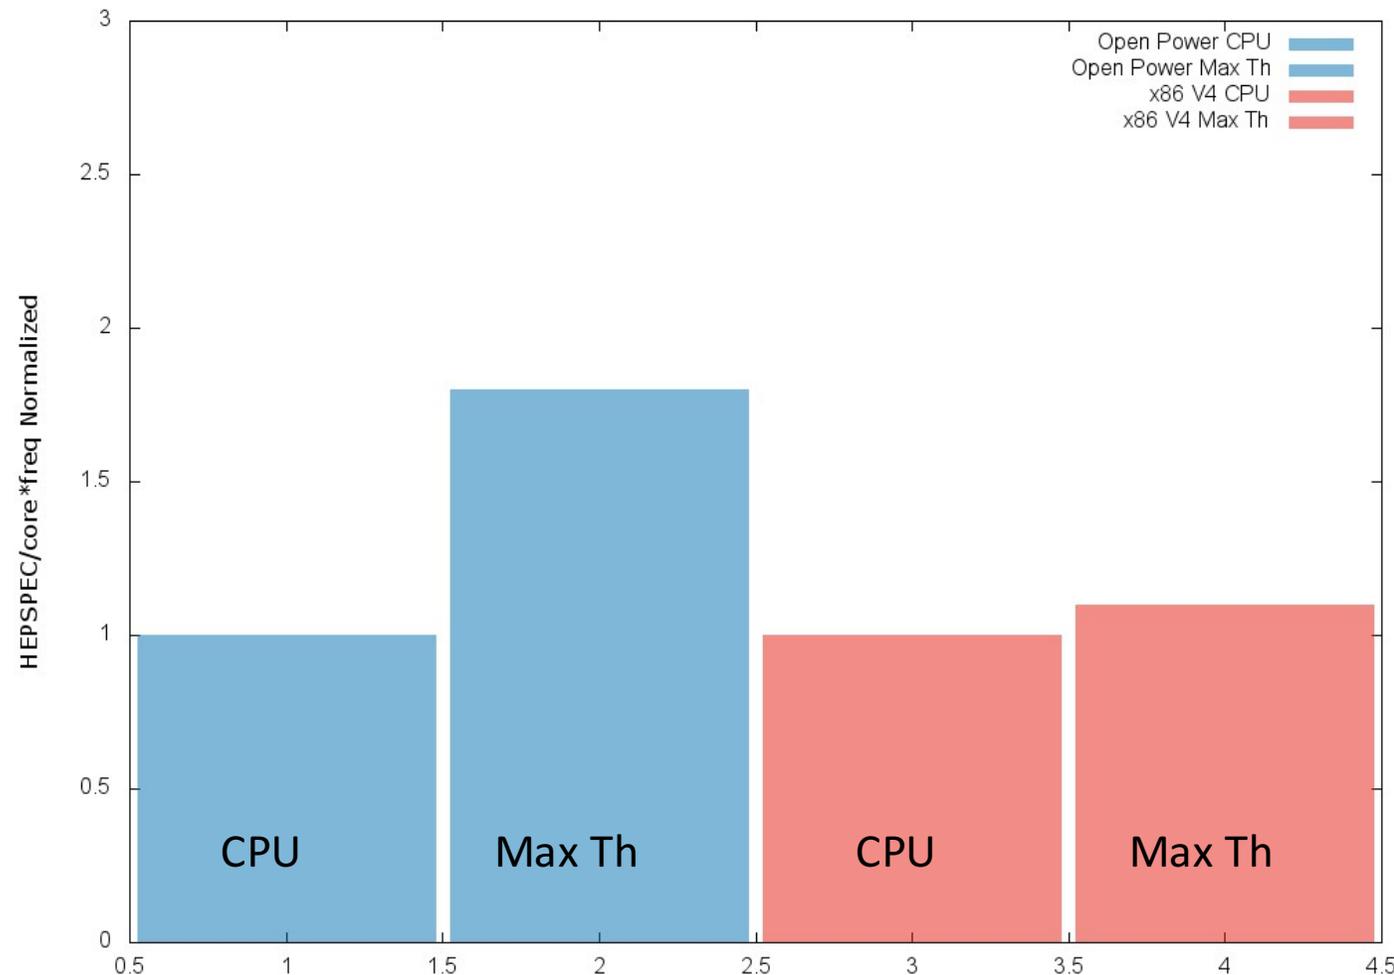
HPC & ENTERPRISE SOLUTIONS



HEPSPEC/(core*freq) per CPU Normalized: Open Power vs x86 V4



HEPSPEC/(core*freq) Normalized: Open Power vs x86 V4



STREAM

HPC & ENTERPRISE SOLUTIONS



Memory performance Benchmark: STREAM <http://www.cs.virginia.edu/stream/> + allocazione dinamica della memoria

COPY: $a(i) = b(i)$

SCALE: $a(i) = q * b(i)$

SUM: $a(i) = b(i) + c(i)$

TRIAD: $a(i) = b(i) + q * c(i)$

OpenPower8 2CPU 8core per CPU @ 3.8GHz, 512 GB RAM

OS: CentOS 7.2.15.11 ppc64le Compiler: gcc 4.8.5

Intel(R) Xeon(R) CPU E5-2697A v4 @ 2.60GHz, 128 GB RAM

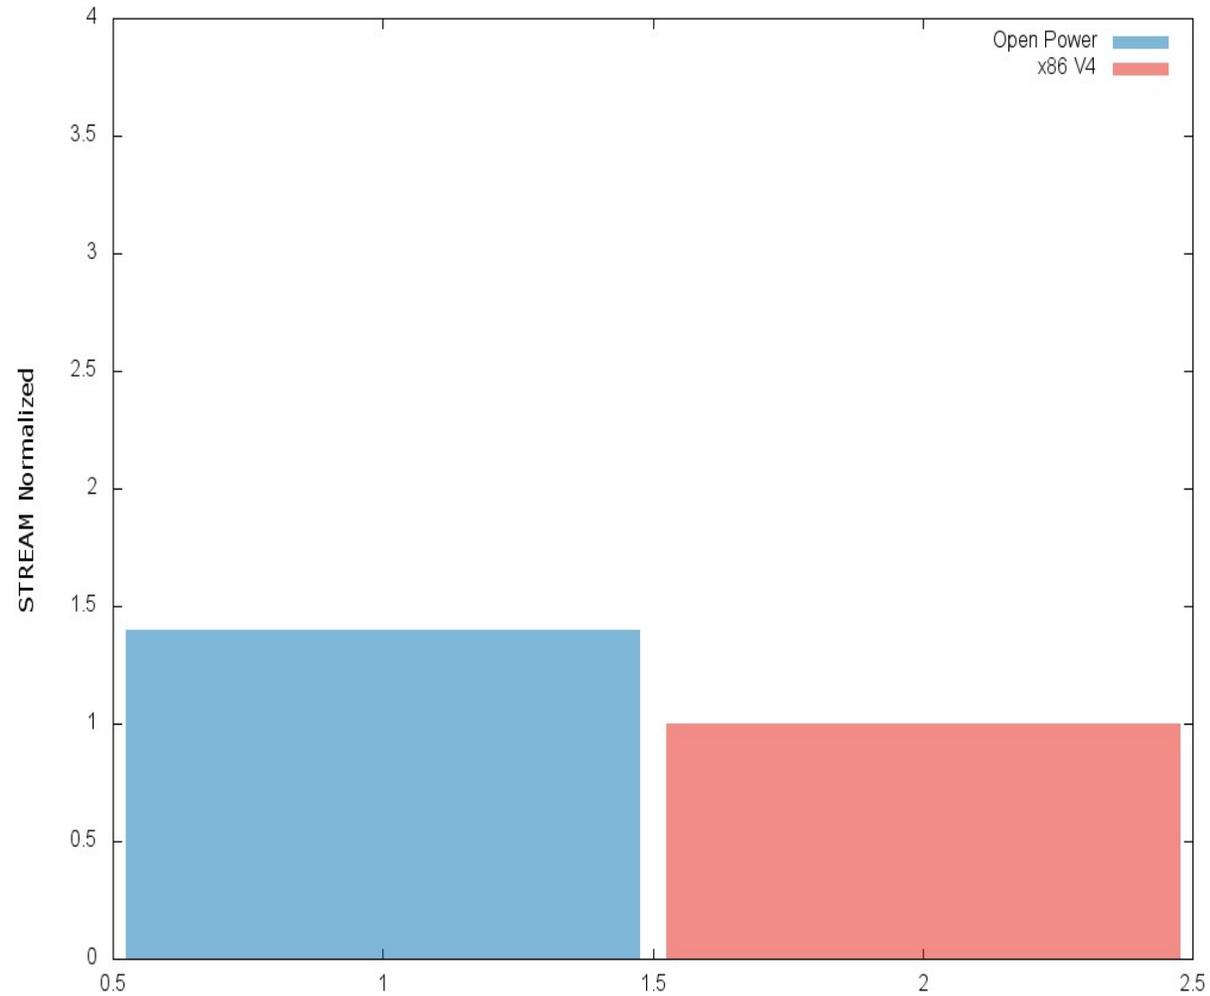
OS: CentOS 6.6 Compiler: gcc 4.4.7

Results Stream

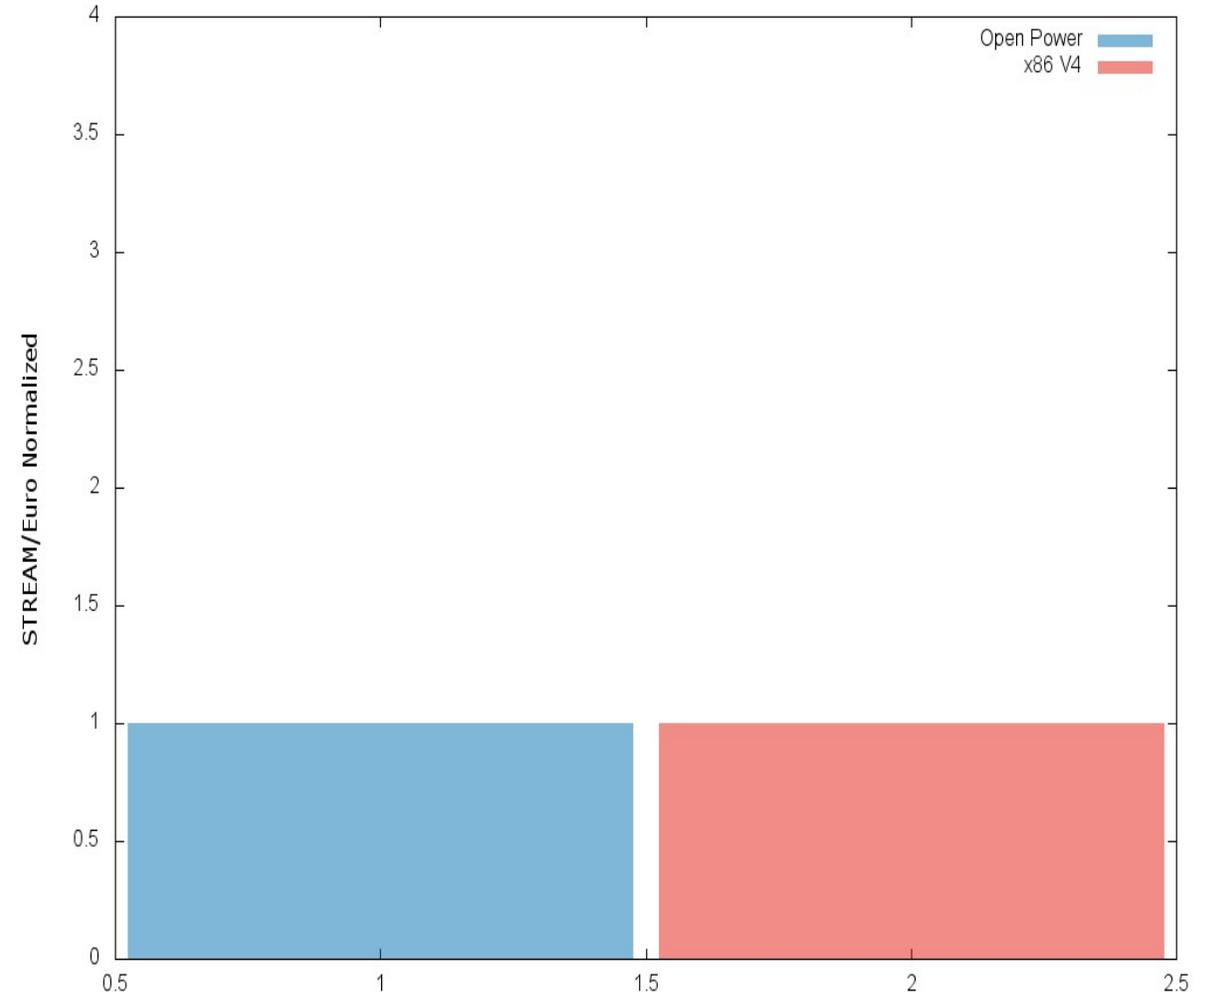
HPC & ENTERPRISE SOLUTIONS

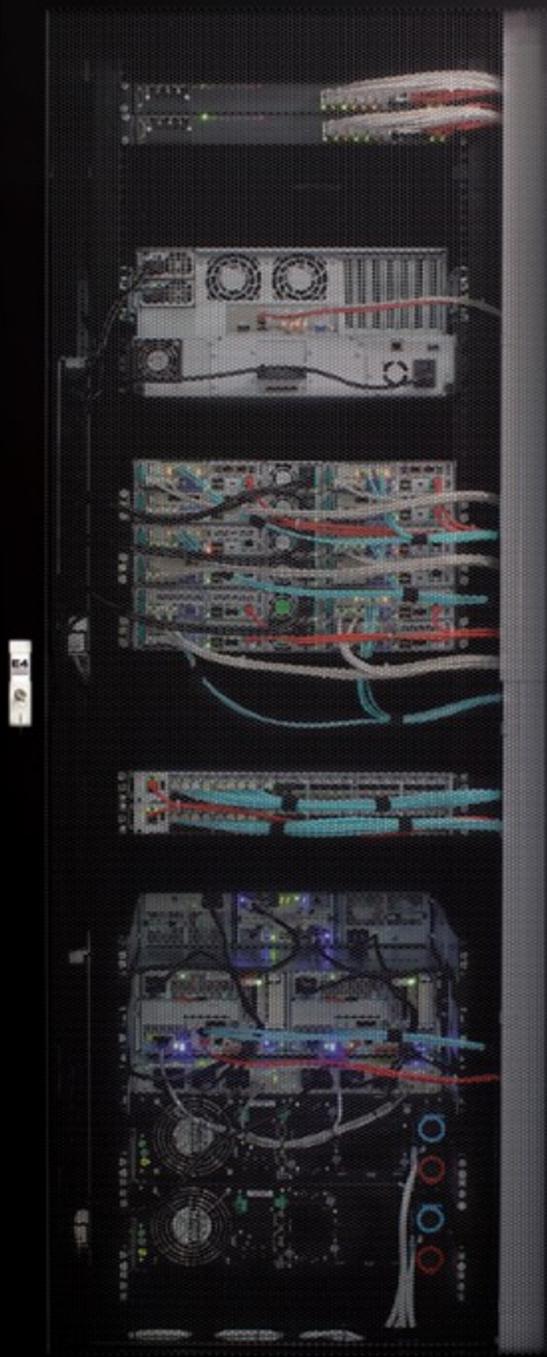


STREAM Normalized: Open Power vs x86 V4



STREAM/Euro Normalized: Open Power vs x86 V4





THANK YOU



E4 Computer Engineering S.p.A.
Via Martiri della libertà, 66
42019 – Scandiano (RE) – Italia

www.e4company.com