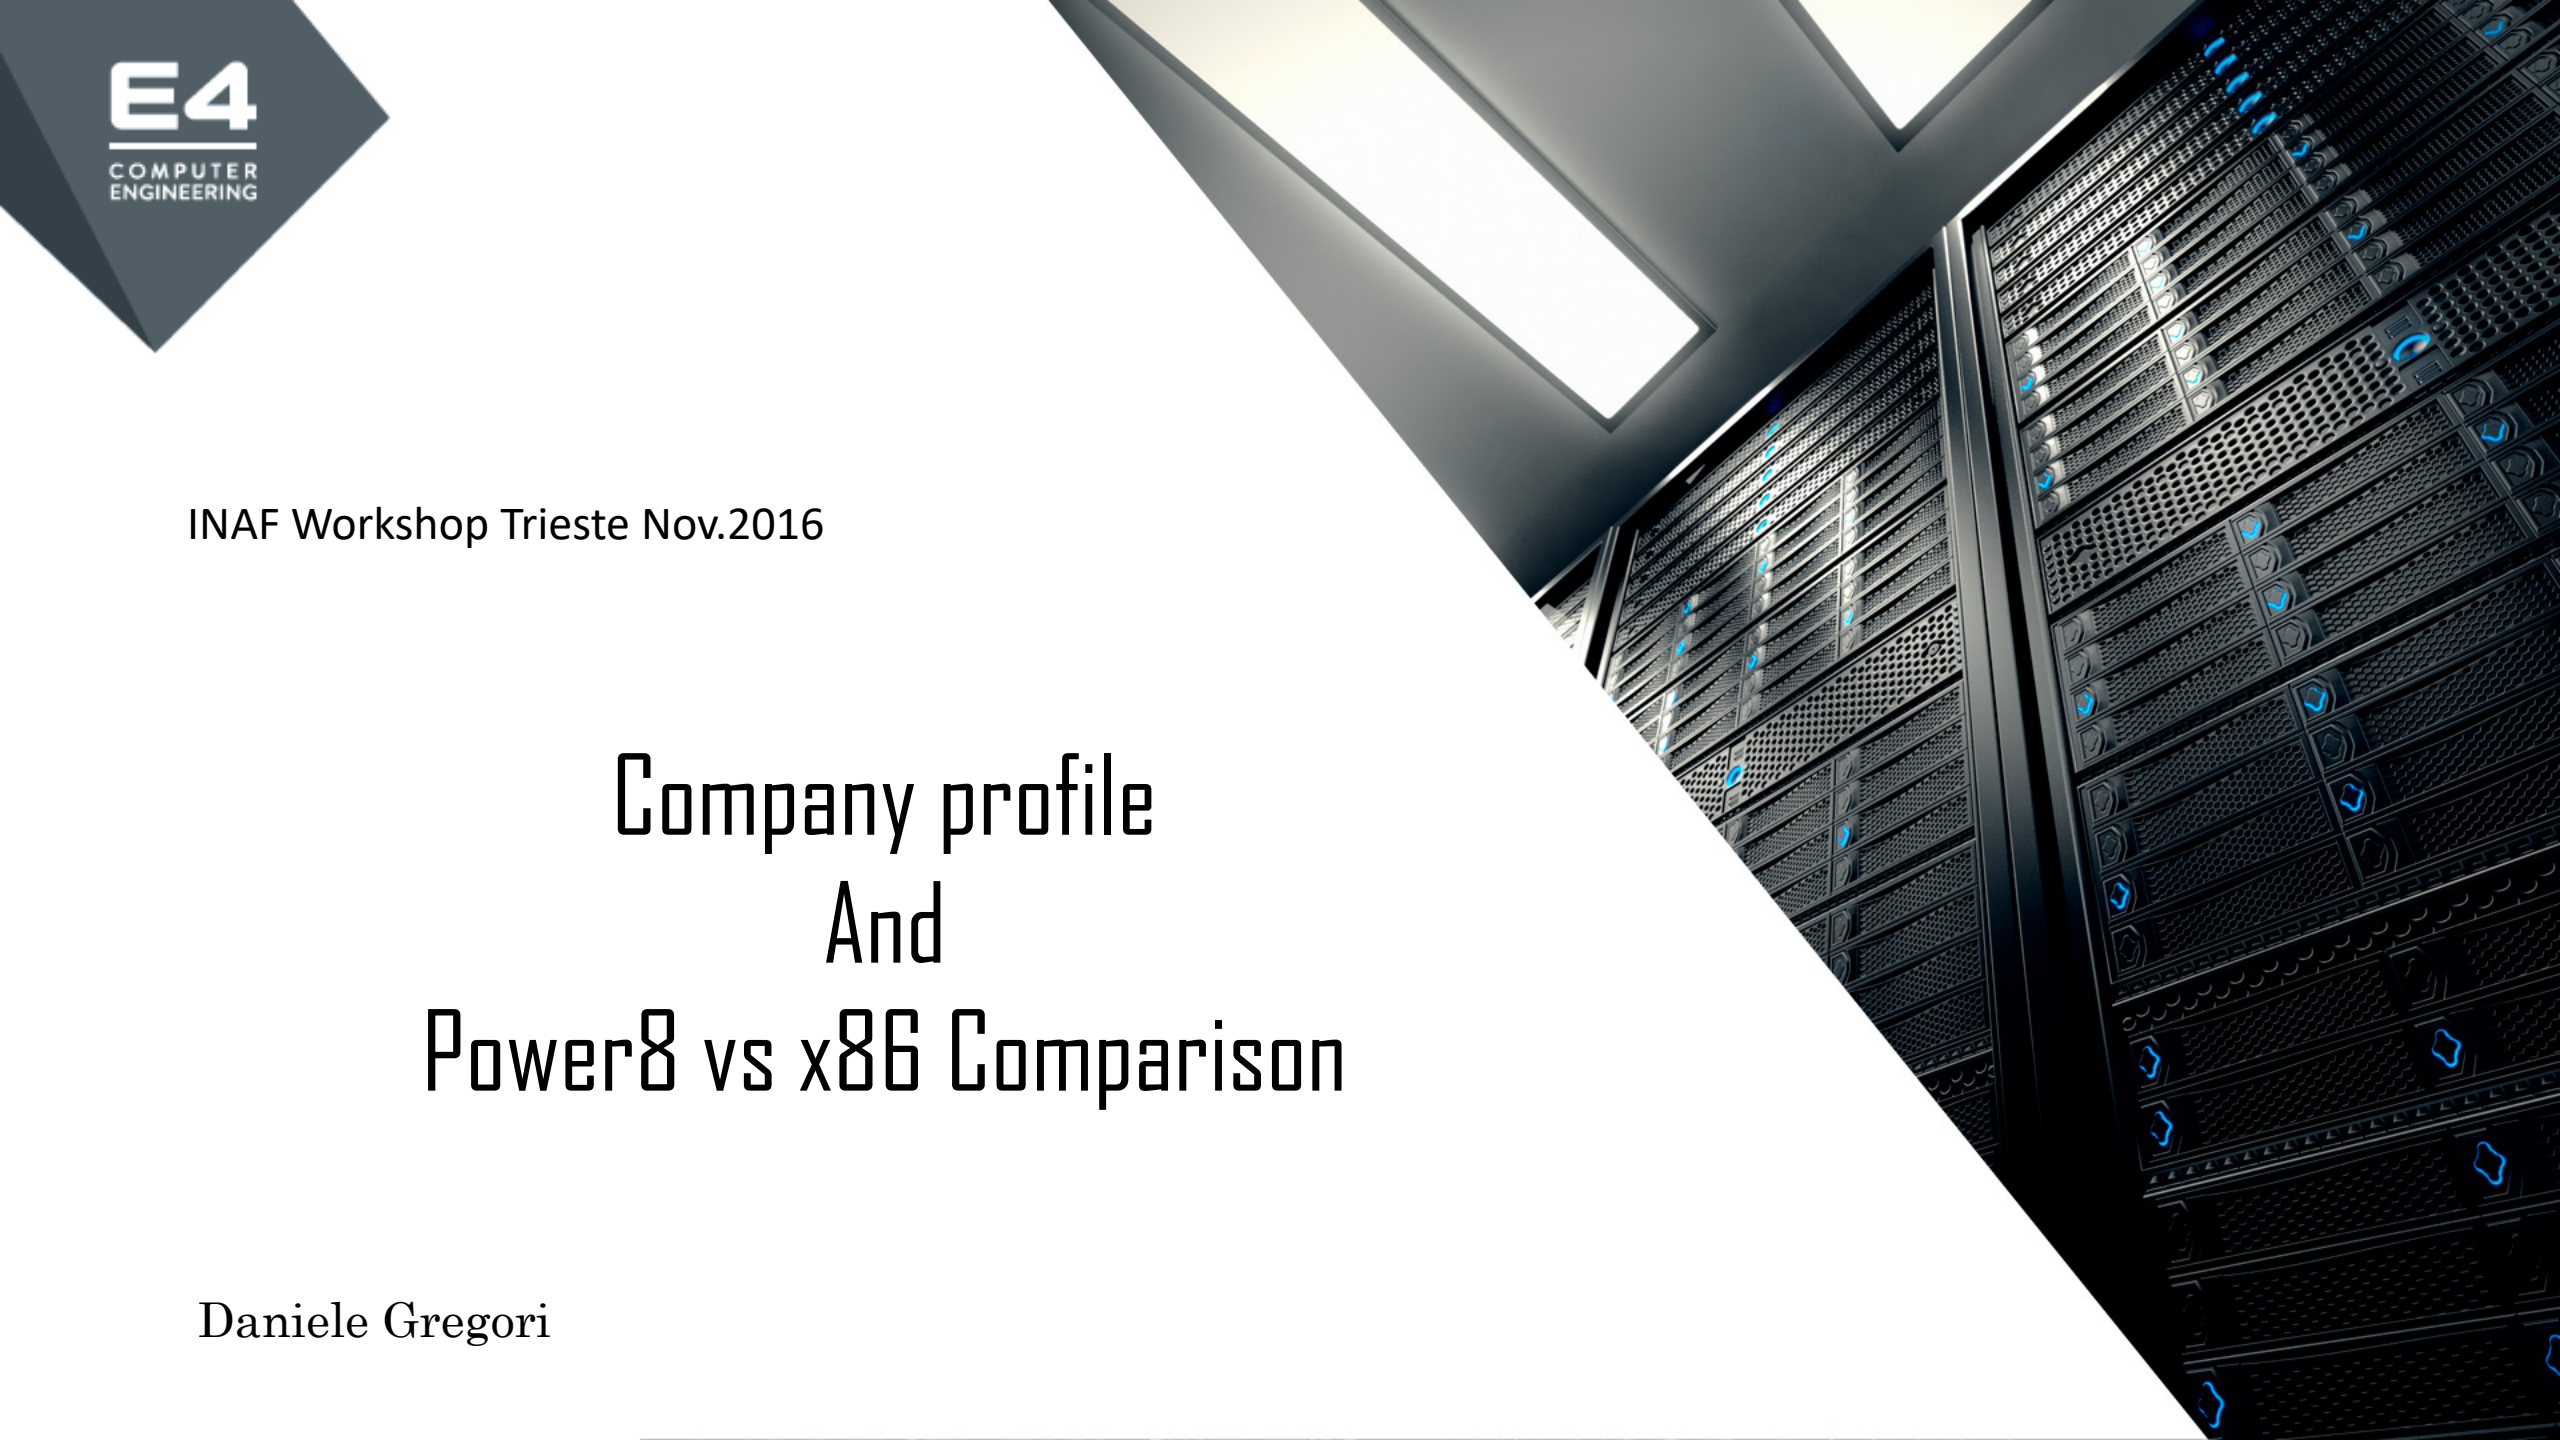

INAF Workshop Trieste Nov.2016

Company profile
And
Power8 vs x86 Comparison

Daniele Gregori

TABLE OF CONTENTS

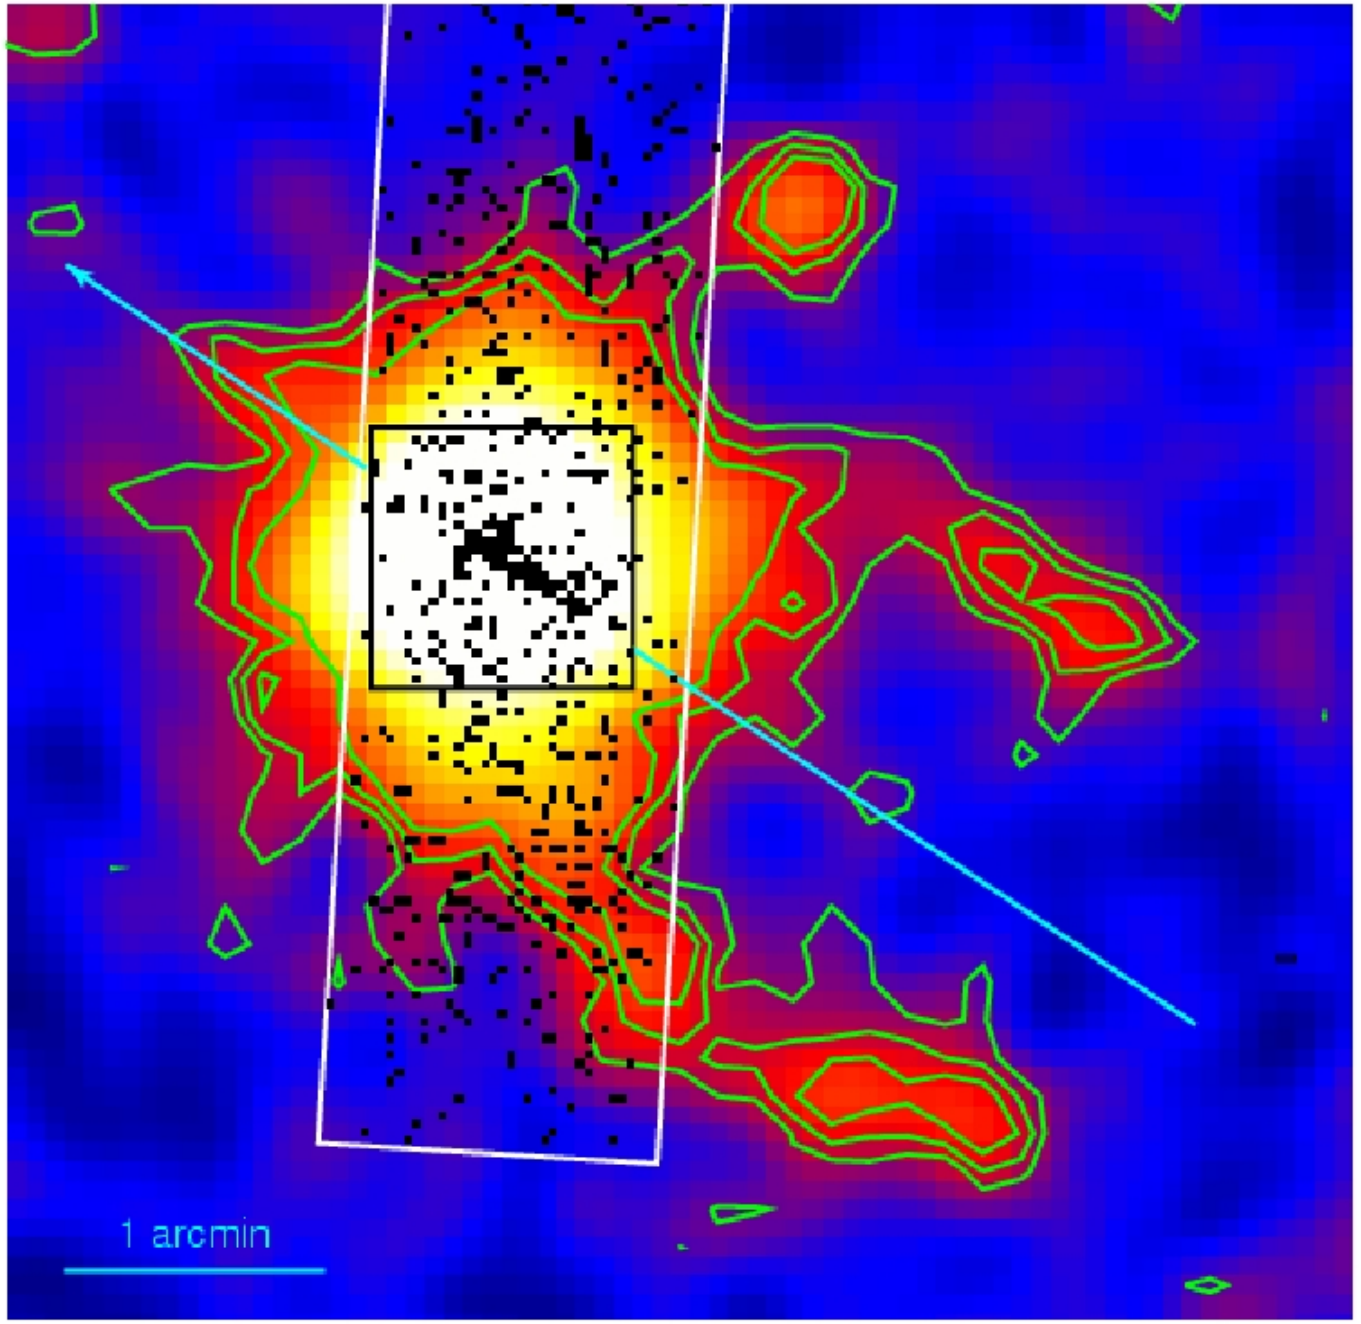

THE COMPANY

Since 2002, E4 Computer Engineering has been innovating and actively encouraging the adoption of new computing and storage technologies. Because new ideas are so important, we invest heavily in research and hence in our future. Thanks to our comprehensive range of hardware, software and services, we are able to offer our customers complete solutions for their most demanding workloads on: HPC, Big-Data, AI, Deep Learning, Data Analytics, Cognitive Computing and for any challenging Storage and Computing requirements.

E4. When Performance Matters.

COMPANY MILESTONES

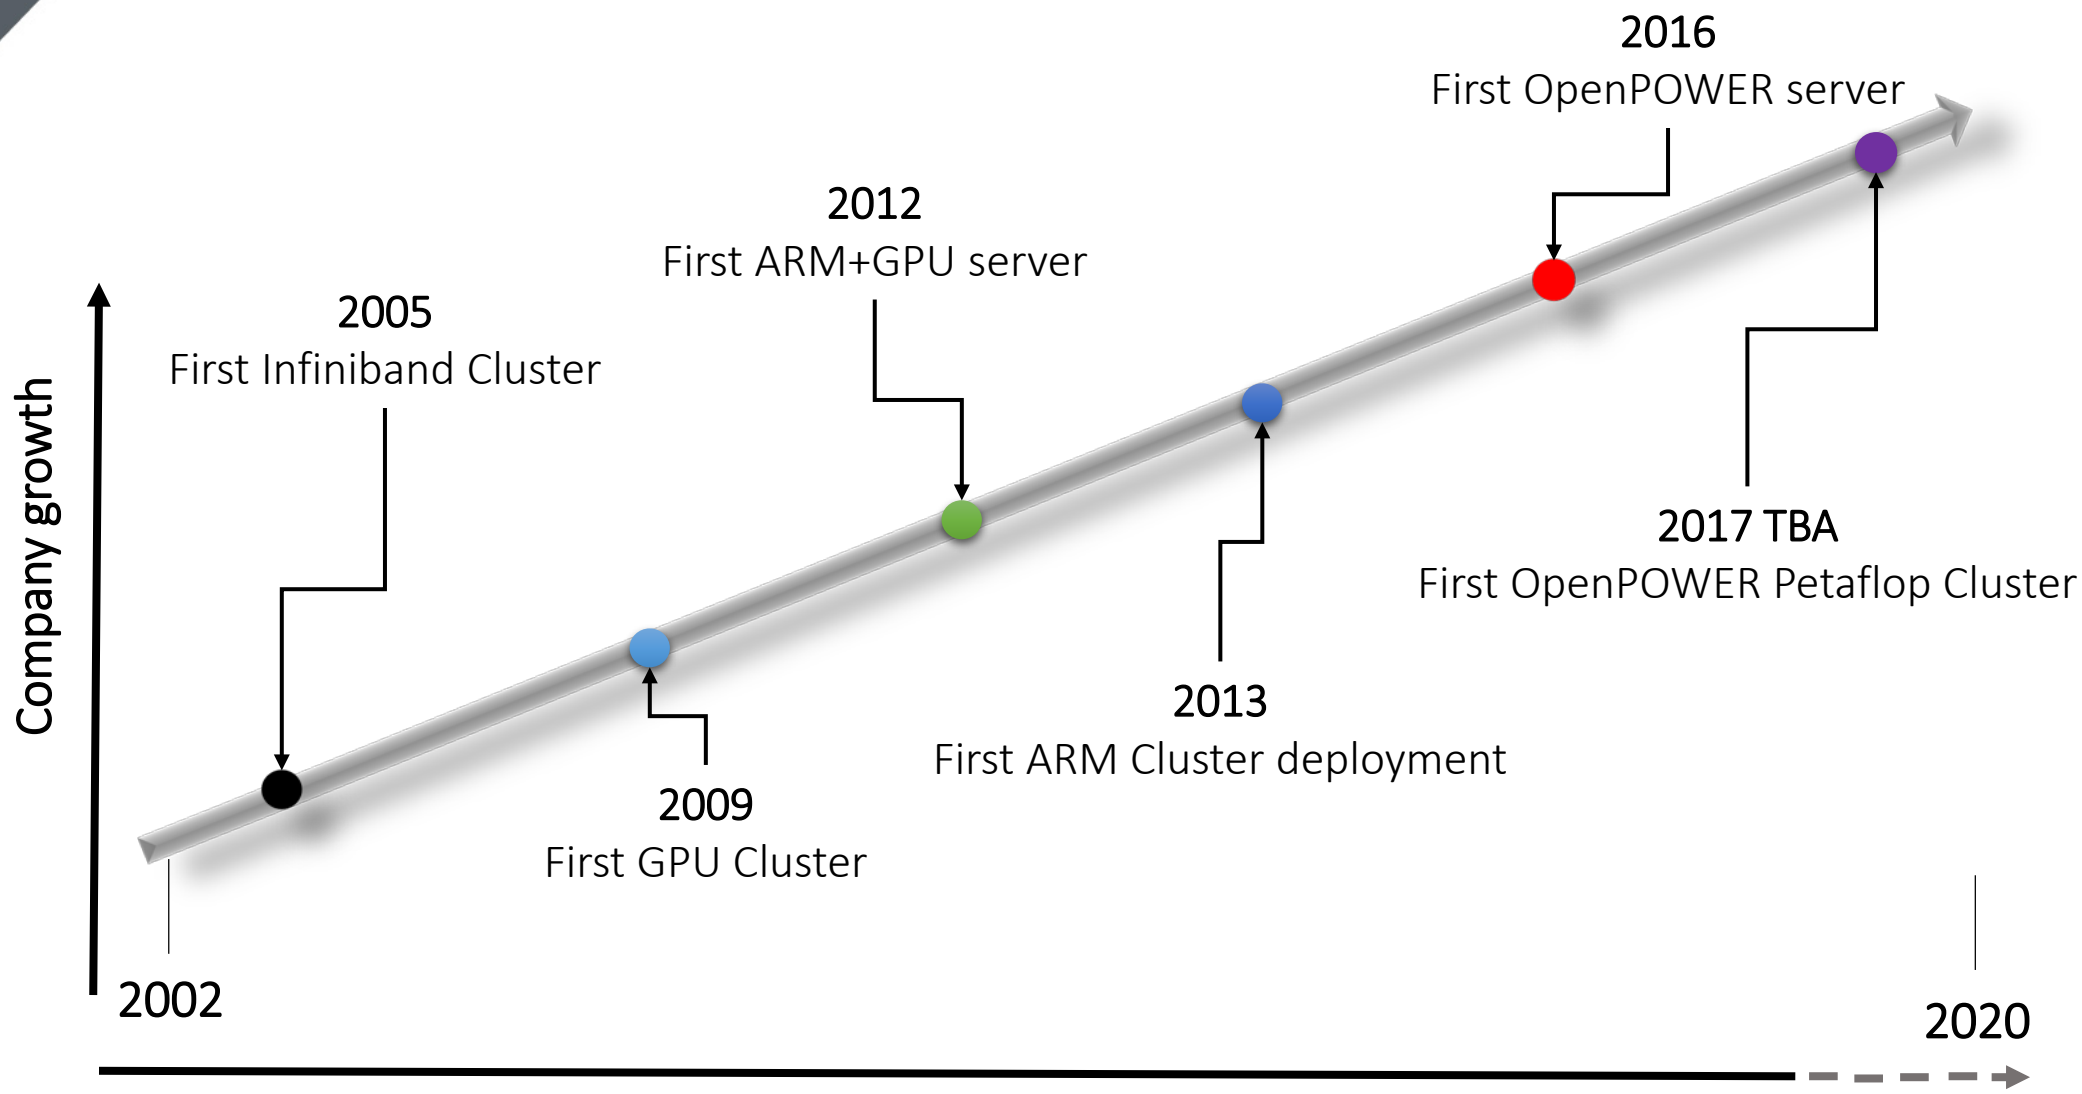

CERN

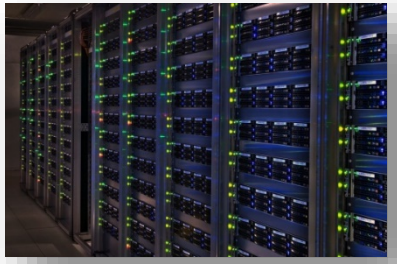

REQUIREMENTS	High density computational nodes Big data storage
CHALLENGES	Delivering standard commodity hardware Providing high performances combined with energy efficiency Ensuring very low failure rate
SOLUTION	5.600+ dual socket mainboards (61.000+ cores) 35.000+ enterprise class hard disks (100PB Storage)
APPLICATION	Grid Computing

INFN

CHALLENGES	Delivering standard commodity hardware Providing combo of high performances & energy efficiency Ensuring very low failure rate
SOLUTION	12PB high performance storage (CNAF) 5PB direct attached storage (Alice – CMS) 4.500 server dual socket (~ 40k computing cores) Several GPU systems 4h intervention times
APPLICATION	Grid Computing

WHAT WE DO

E4 COMPANY PILLARS

HARDWARE PRODUCTS

EXTREME COMPUTING SOLUTIONS

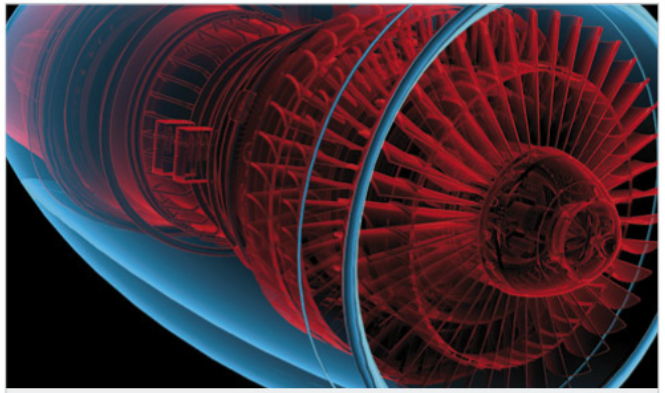

R&D PROTOTYPING

HIGH AUTOMATION FOR PRODUCTION

Our production takes place through high automation tools for installation, configuration and quality control.

I.E. over 2016 there have been two major CERN productions where we have manufactured ~ 2300 servers and 200 JBOD.

The software tools adopted are Open Source: Cobbler – Puppet – Nagios – Bugzilla – Wiki.
We internally support the “E4tools” application development for quality control.

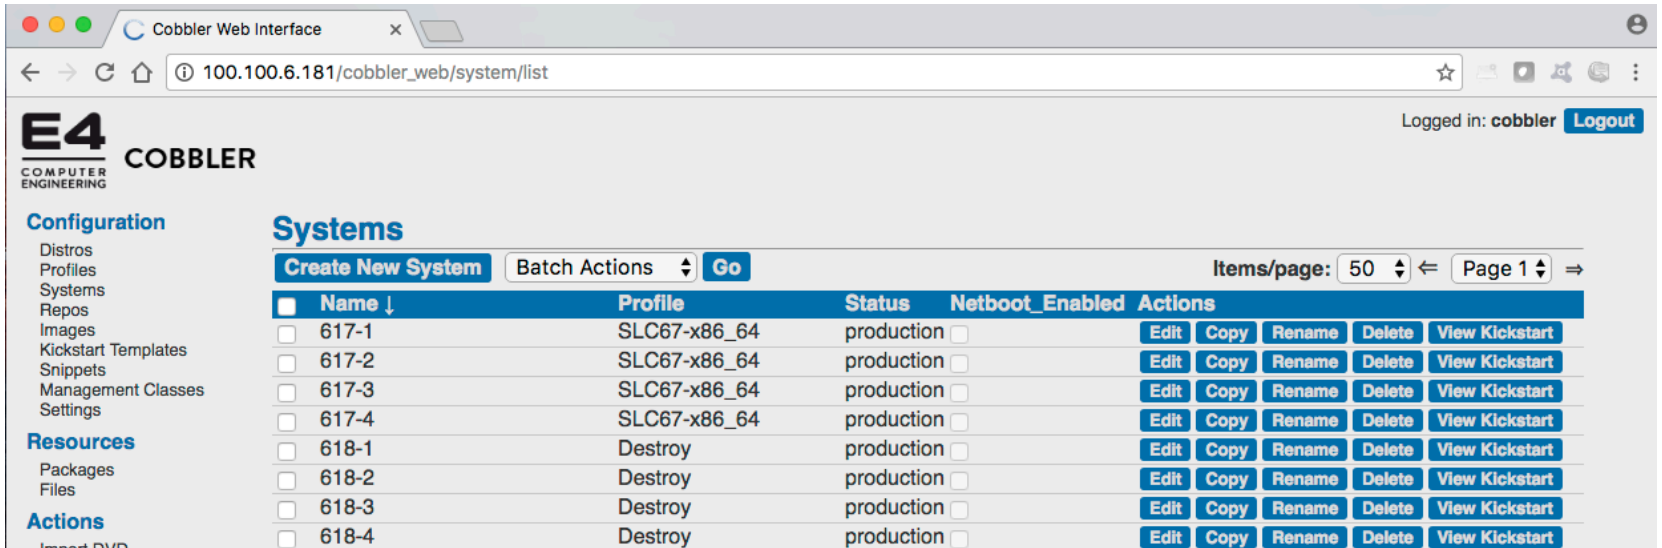

Server Quality

E4 TOOLS

A PROPRIETARY SOFTWARE FOR SYSTEMS' SET-UP AUTOMATION AND HOMOGENEITY CHECKS

- FIRMWARES' CHECKS, UPGRADE, CONFIGURATION
- SETUP CHECKS
- HOMOGENEITY CHECKS
- SANITY CHECKS
- PERFORMANCE CHECKS
- DETAILED REPORTS
- HARDWARE CATALOG, PERFORMANCE REPORT, ERROR REPORT
- SUMMARY REPORT FOR EACH MACHINE PRODUCED

New R&D LAB

- 30 m^2
- temperature 27/30°c
- 6 x Rack 19'
- 4 x Chiller 22 kw
- Active Power available ~100 kw
- Gestione Hardware via OpenDCIM open source SW

R&D CURRENT ACTIVITIES

LU·M·E·A
Lustre Appliance Made Easy

OpenPower
HPC cluster

ARM
HPC cluster

SOLUTIONS DEVELOPMENT

HPC Opensuite 2.0

Long Term Data Preservation
From bit to software environment protection

RESEARCH
analysis, technologies, measurements...

IEEL
Plug-in for DotHill storages

E4TOOLS
Tools for automation of production and R&D assistance

SOFTWARE DEVELOPMENT

E4
COMPUTER
ENGINEERING

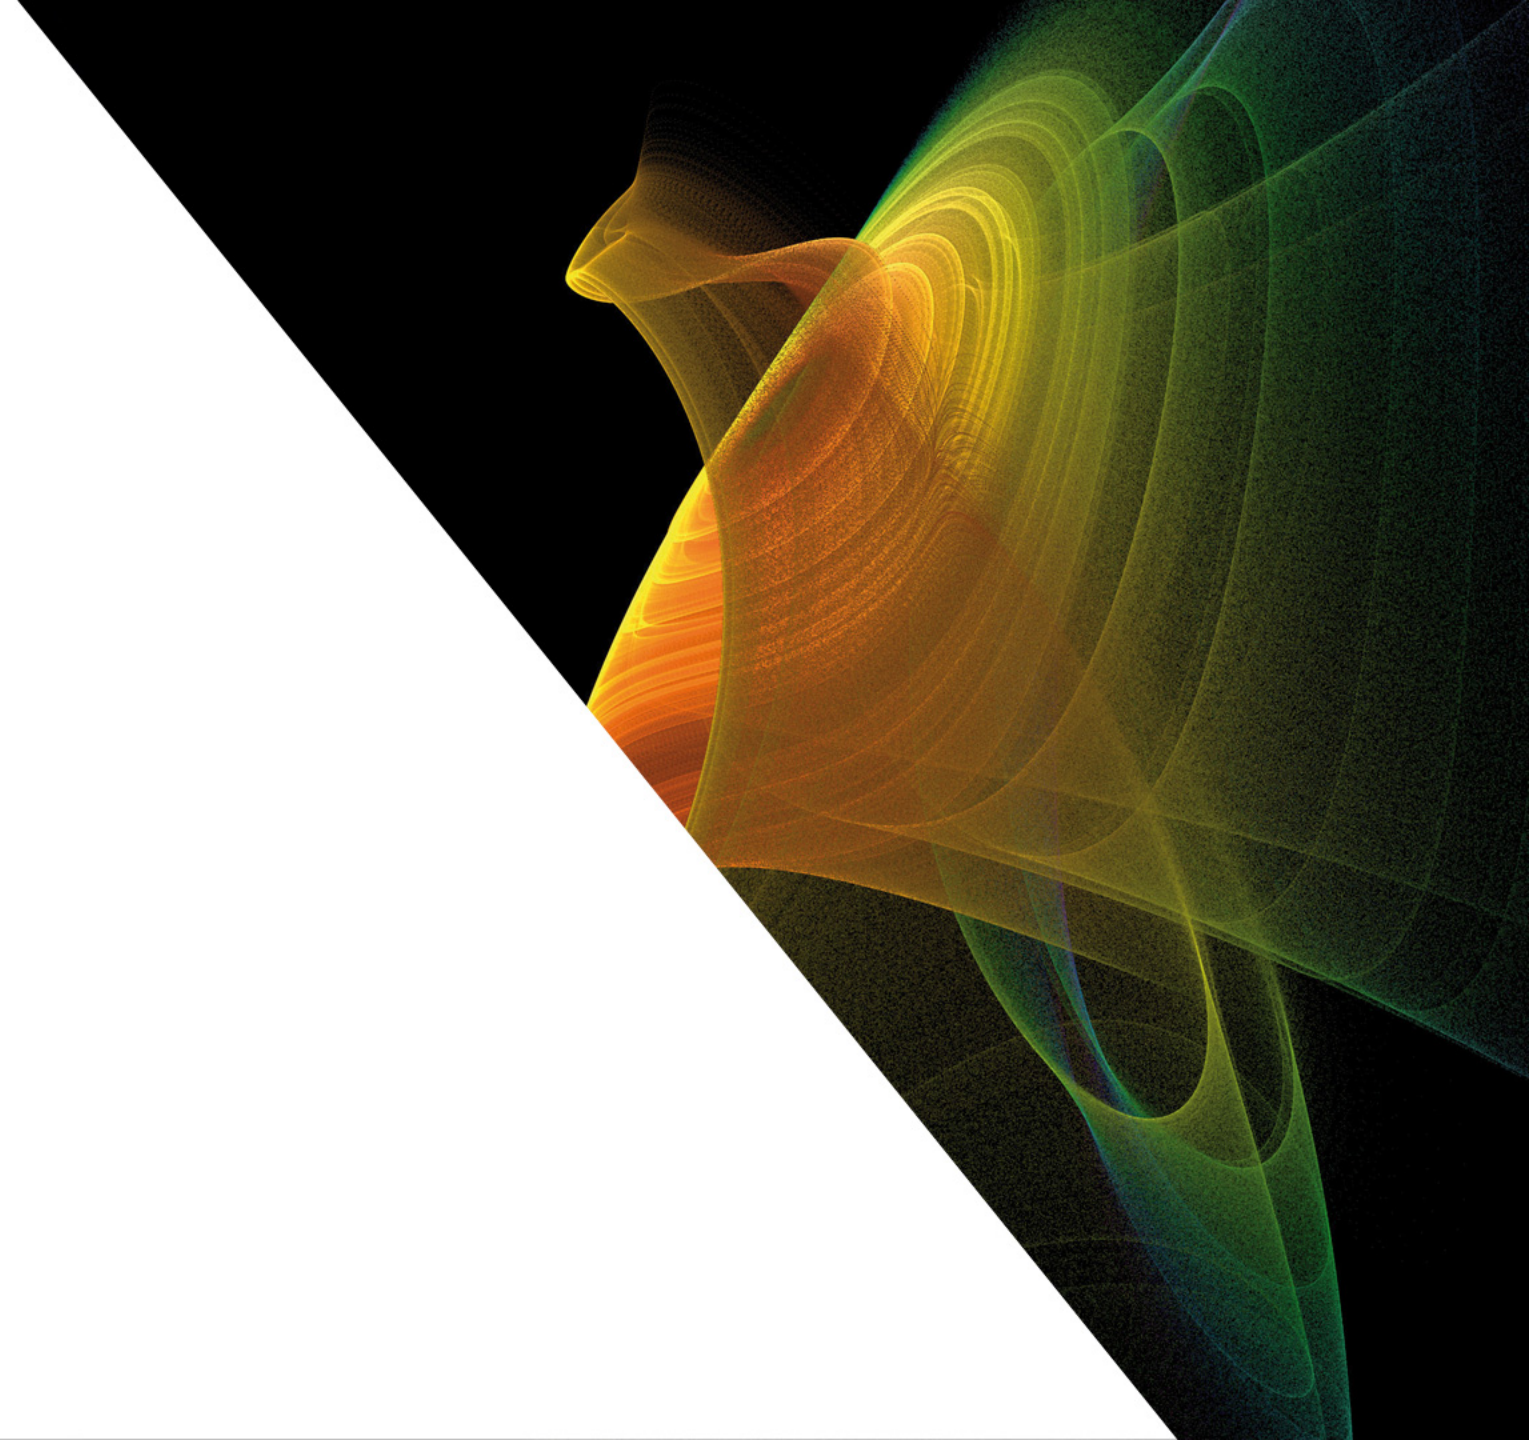

OUR SOLUTIONS

STORAGE & SOFTWARE DEFINED STORAGE

lustre

IBM Spectrum Scale

SPECTRA

DDN STORAGE

panasas

We are able to provide solutions with Parallel Filesystem and hierarchical storage on tape library

HPC Open Suite

E4 HPC Open Suite - THE PERFECT CLUSTER RECIPE

HPC clusters coupled with different management suites, starting from a collection of open source tools up to enterprise management solutions. E4 HPC clusters are supplied by default with E4's HPC OPEN SUITE

NODE PROVISIONING
INSTALLATION
CONFIGURATIONBARE
METAL REINSTALL

NETWORK SETUP
NETWORK SERVICES
STORAGE

USER MANAGEMENT
RESOURCE/ QUEUE
MANAGER

IN-BAND AND
OUT-OF-BAND
MANAGEMENT

CLUSTER COMMAND TOOLS
MONITORING AND ALERTING
POWER MANAGEMENT

DEVELOPMENT TOOLS
DEBUGGING /
PROFILING TOOLS

BENCHMARKING TOOLS
CERTIFICATIONS

Server Open Power8

OpenPOWER SERVER OP205

FIRST POWER8 PROCESSOR-BASED SYSTEM IN ITALY AND EUROPE

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

BENEFITS

- Full stack of tools for cluster management
- Cluster available for remote access
- Performance tuning services

Forthcoming Open Power 8+

OpenPOWER SERVER OP206

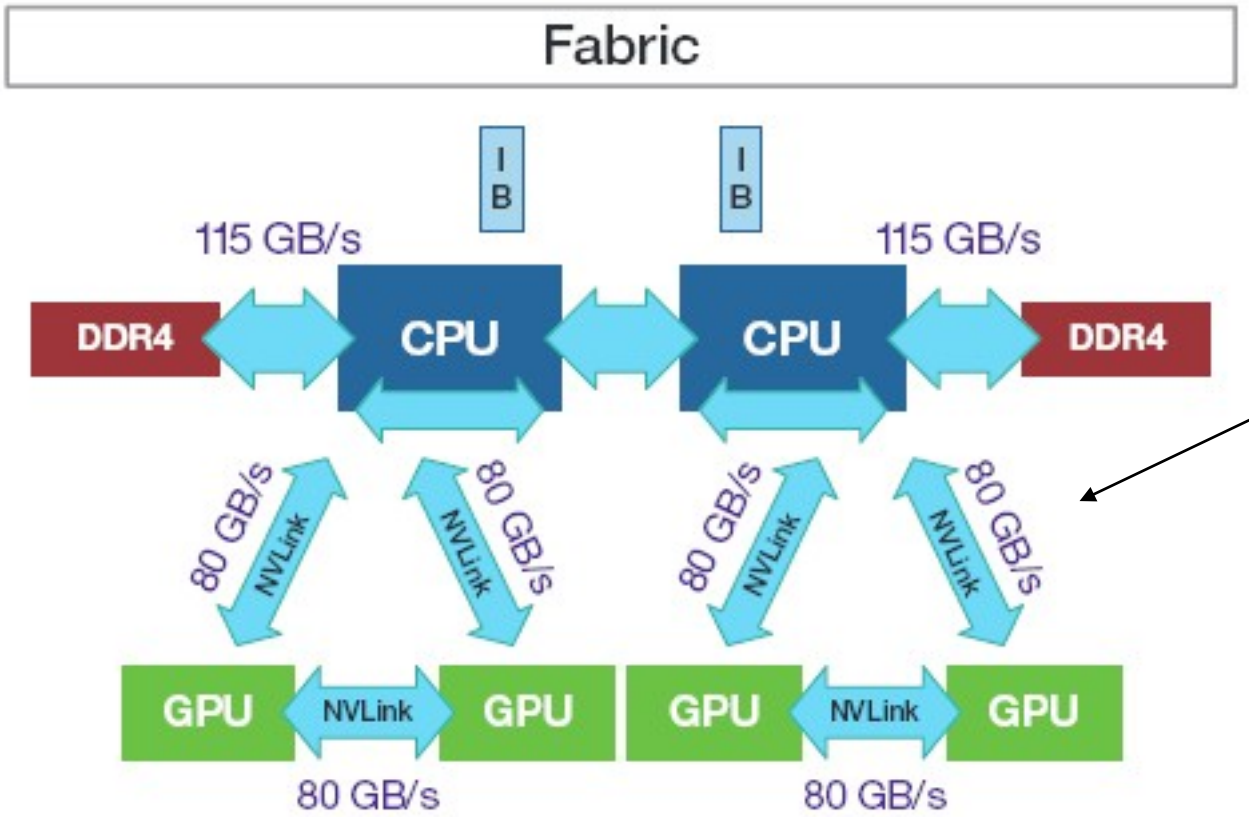

GPGPU-CPU NVLINK
Available only for
Power Architecture

HEPSPEC TEST

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) 2CPU 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.hepix.org/benchmarks/doku.php
https://www.spec.org/

HEP SPEC

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

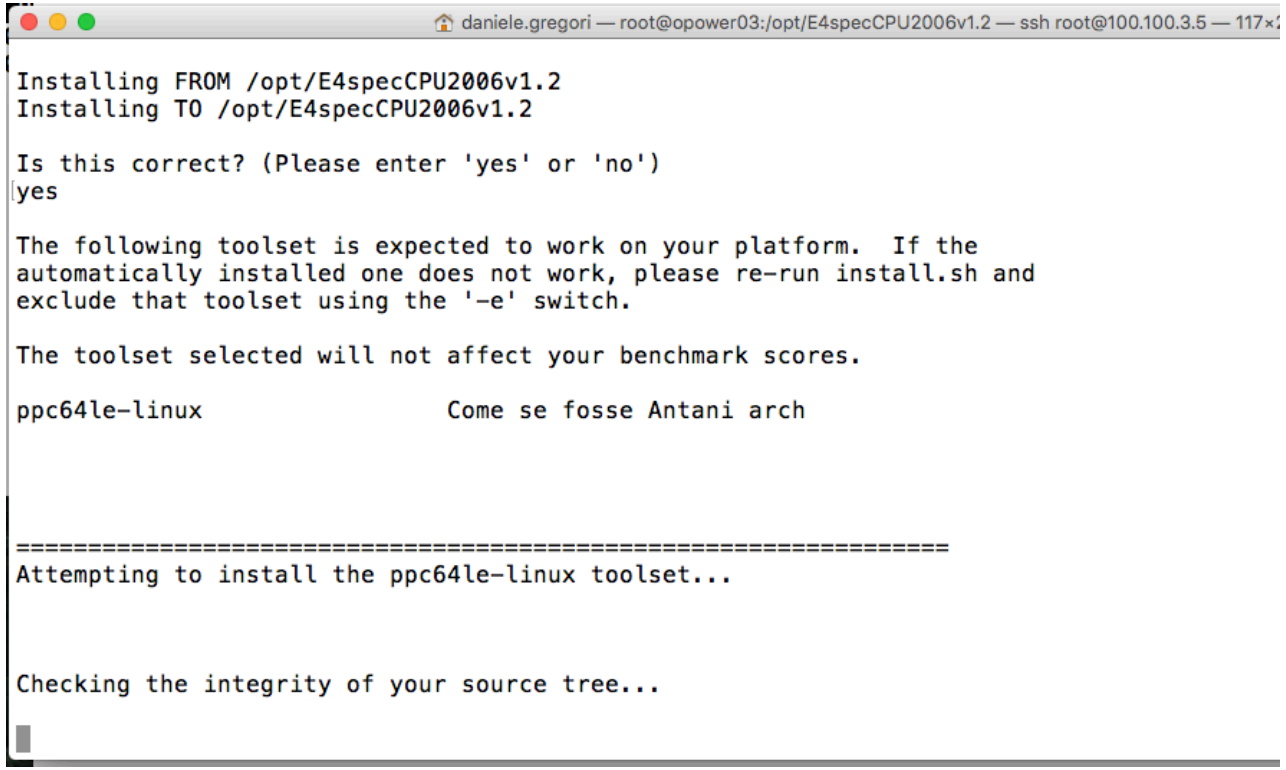

HEPSPEC RESULTS

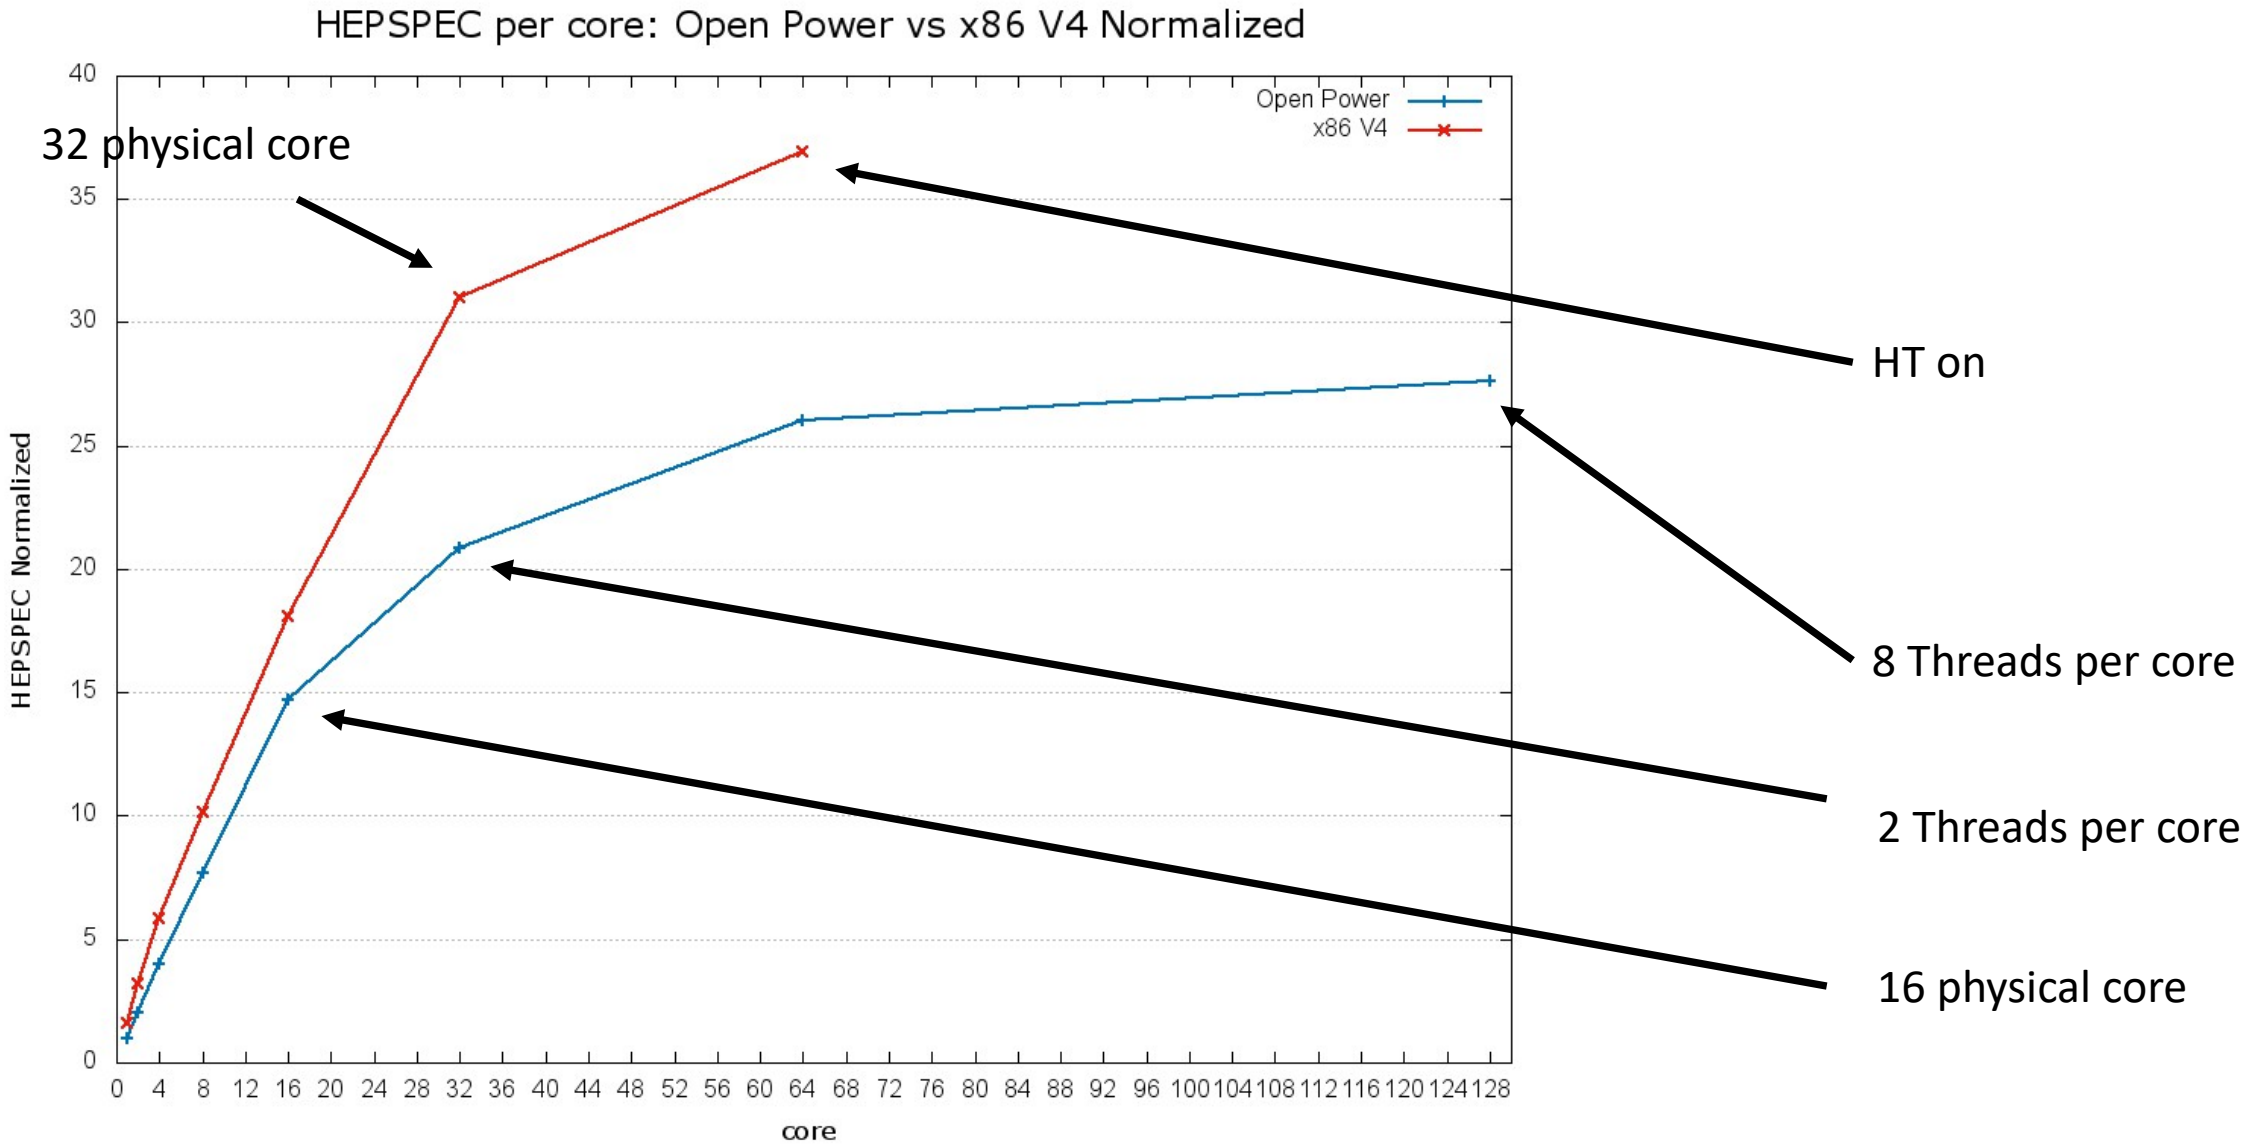

32 physical core

HT on

8 Threads per core

2 Threads per core

16 physical core

HEPSPEC RESULTS

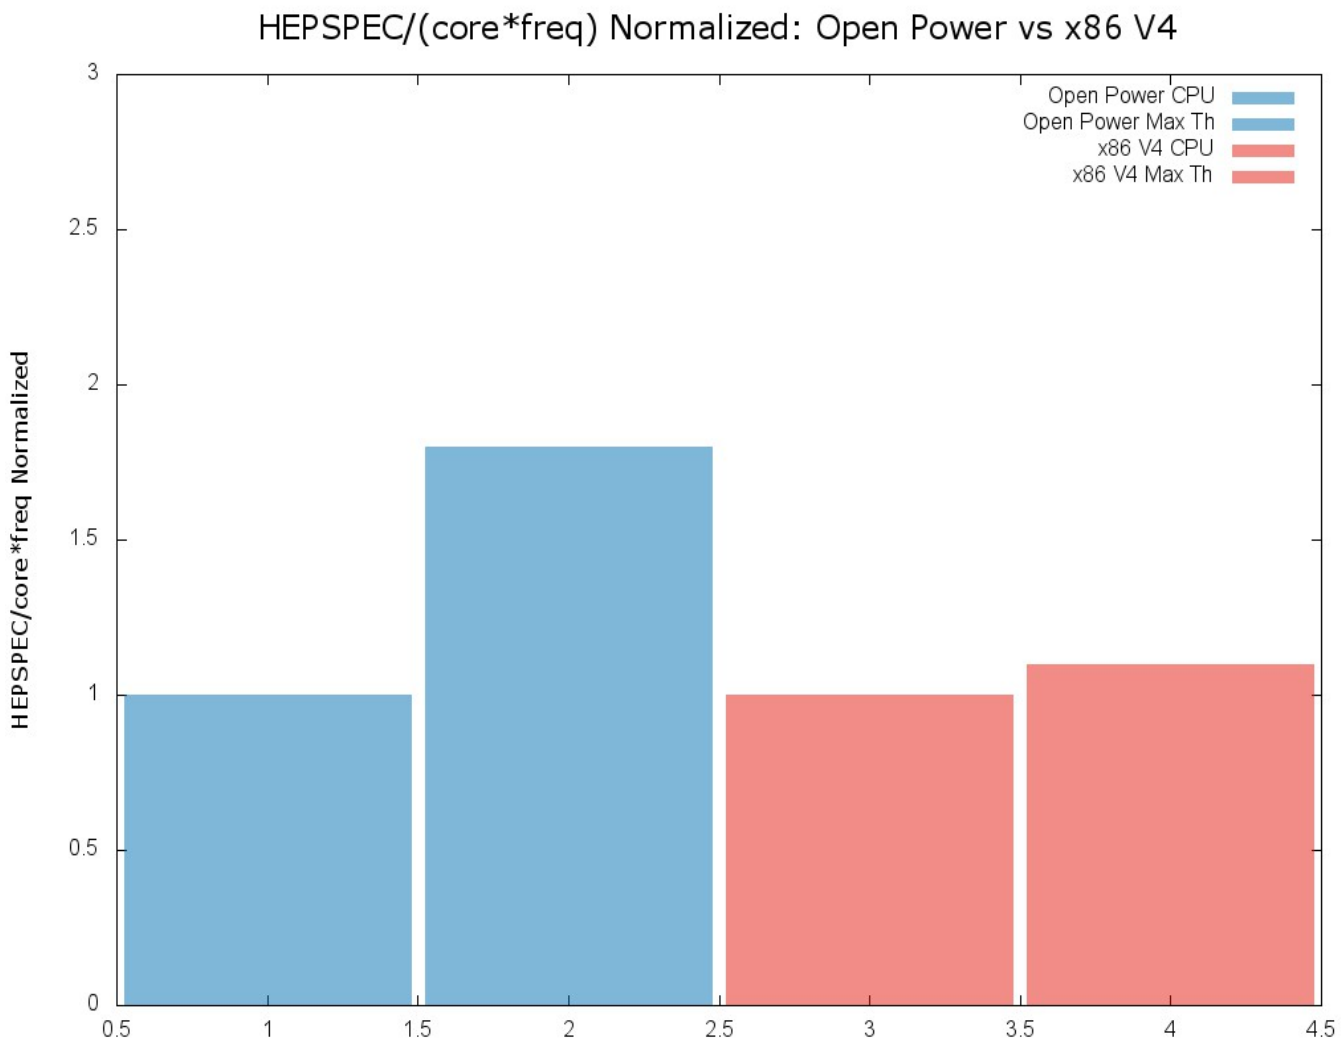

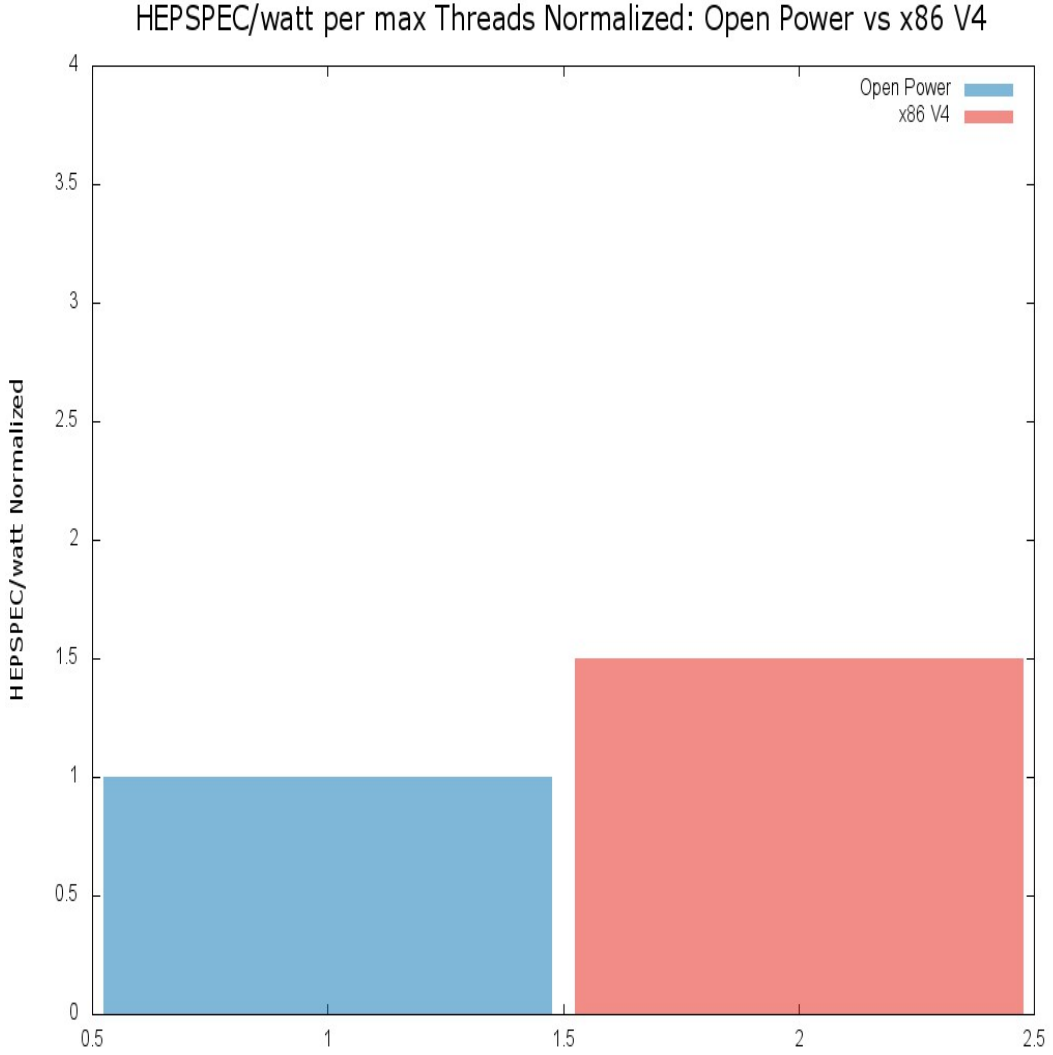

HPC Challenge TEST

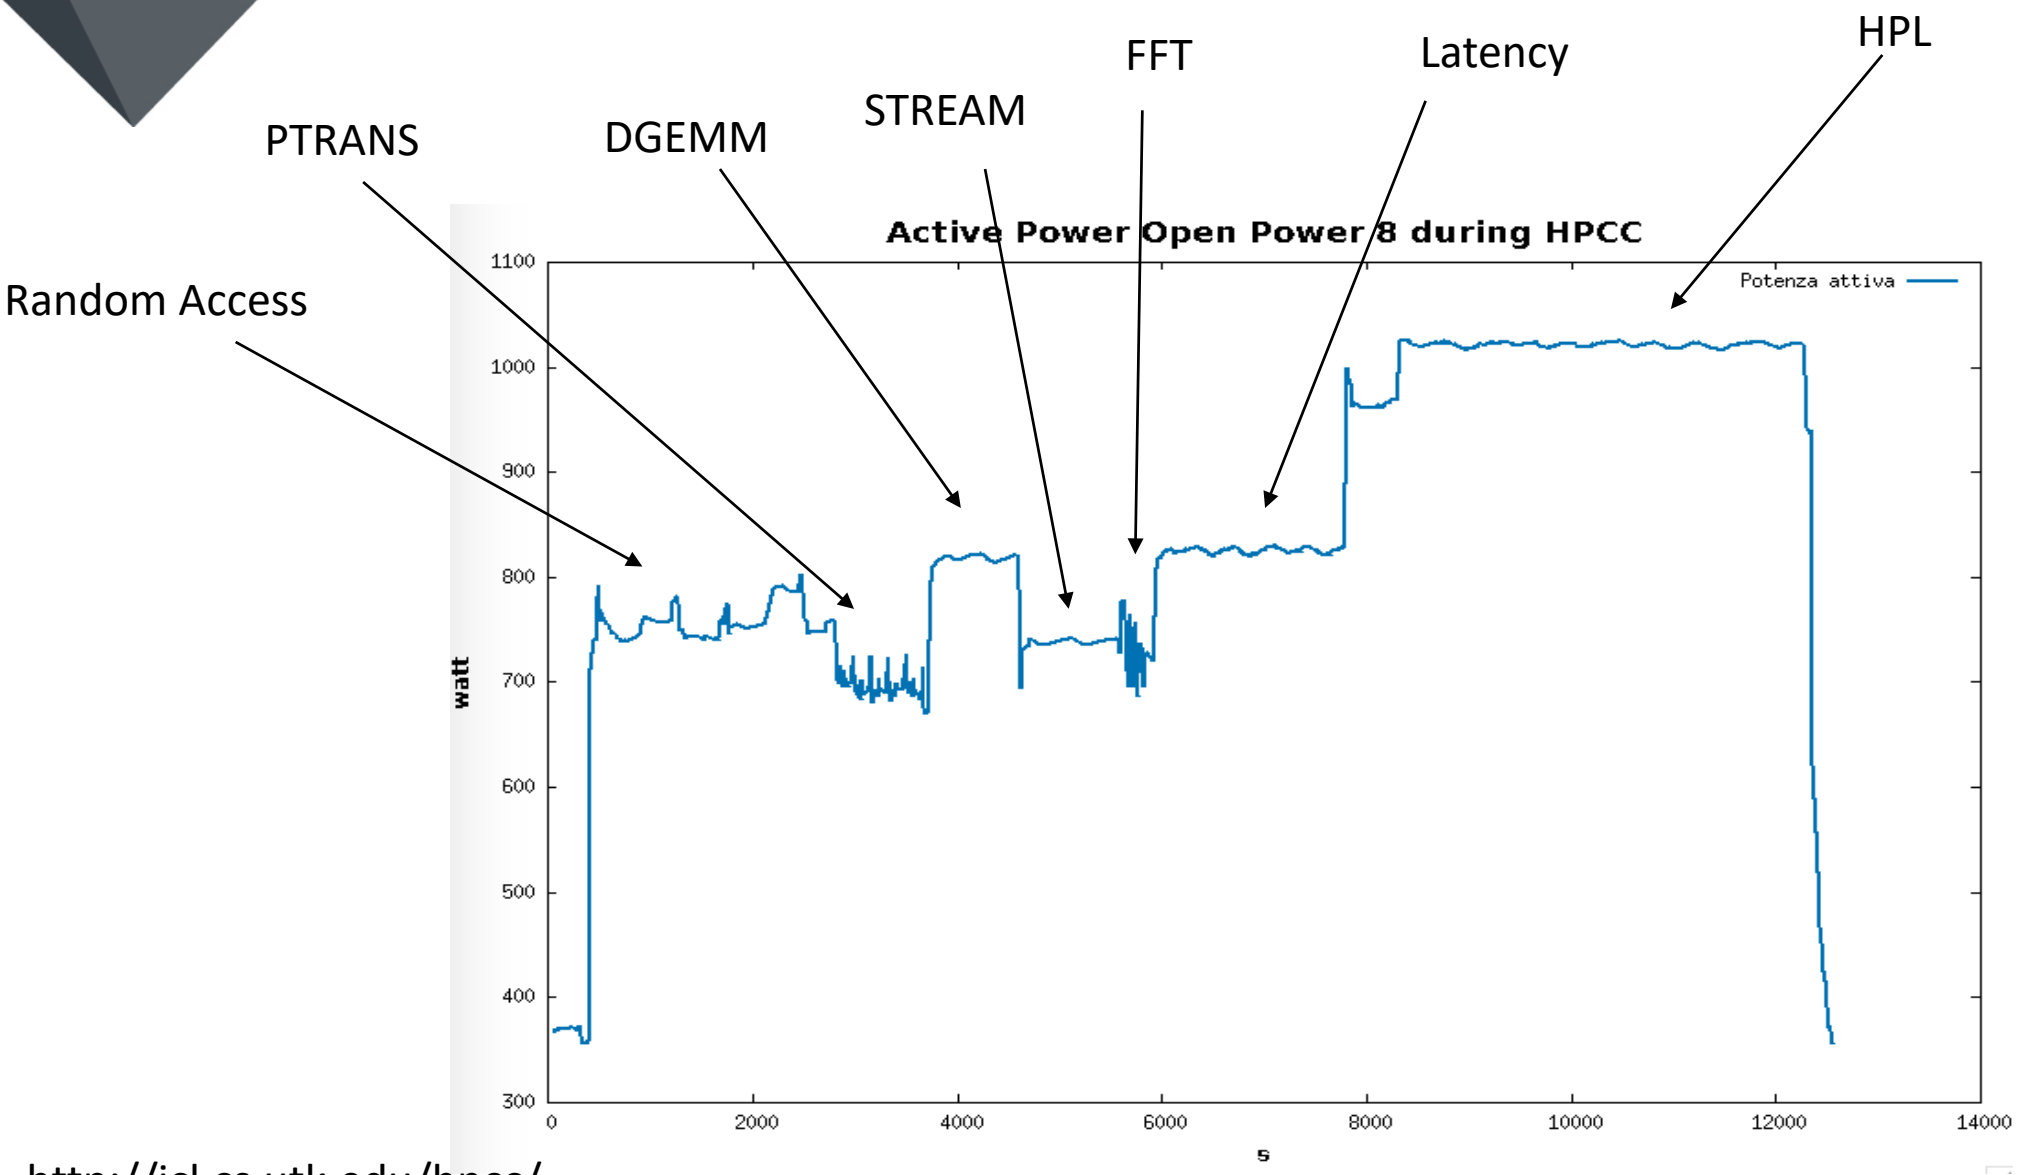

http://icl.cs.utk.edu/hpcc/

STREAM TEST

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

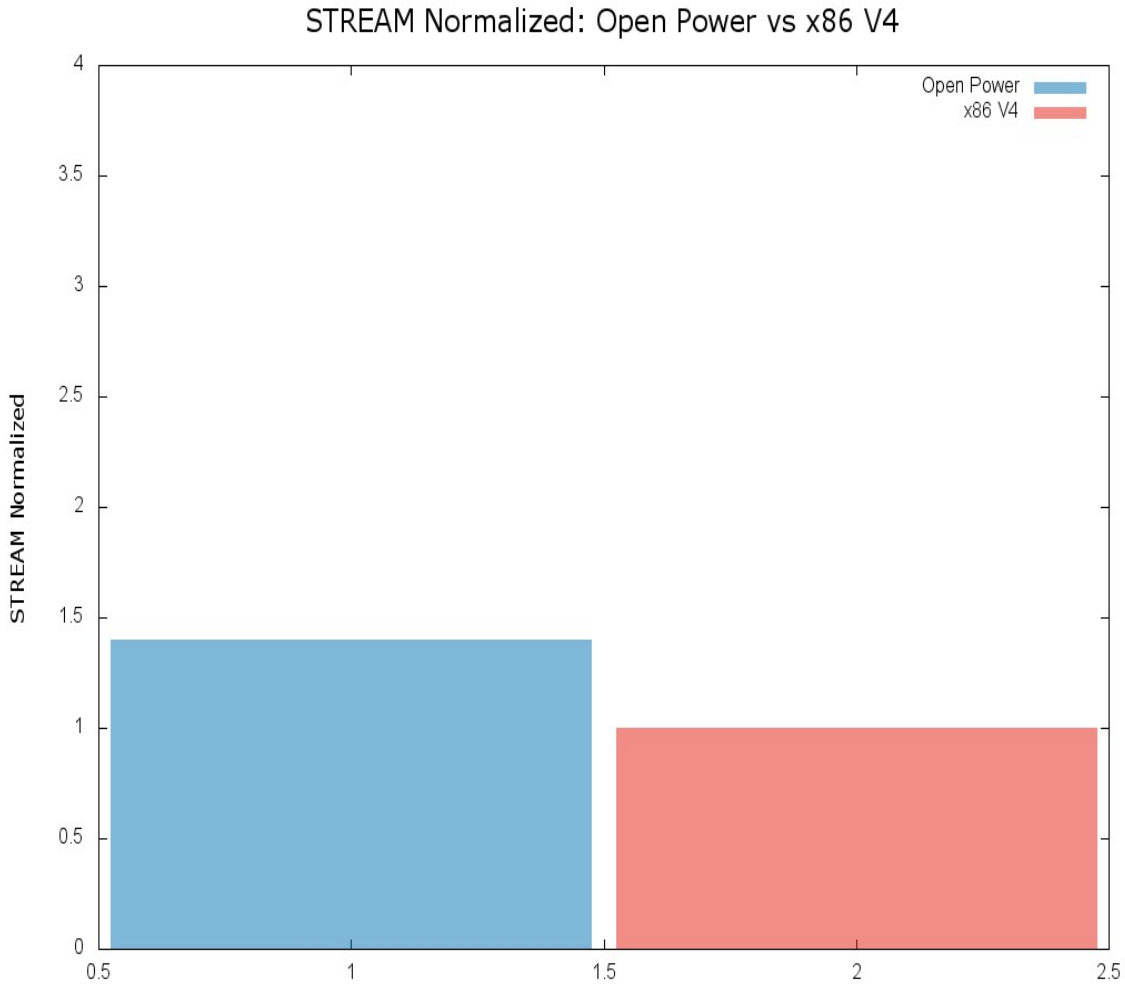

Linpack TEST

http://www.netlib.org/benchmark/hpl/

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

OS: CentOS 7.2.1511 ppc64le Compiler: gcc 4.8.5

Intel(R) Xeon(R) 2CPU 10core per CPU E5-2630v4 @ 2.2GHz, 256 GB RAM

OS: CentOS 7.2.1511 Compiler: Intel Parallel Studio 2017

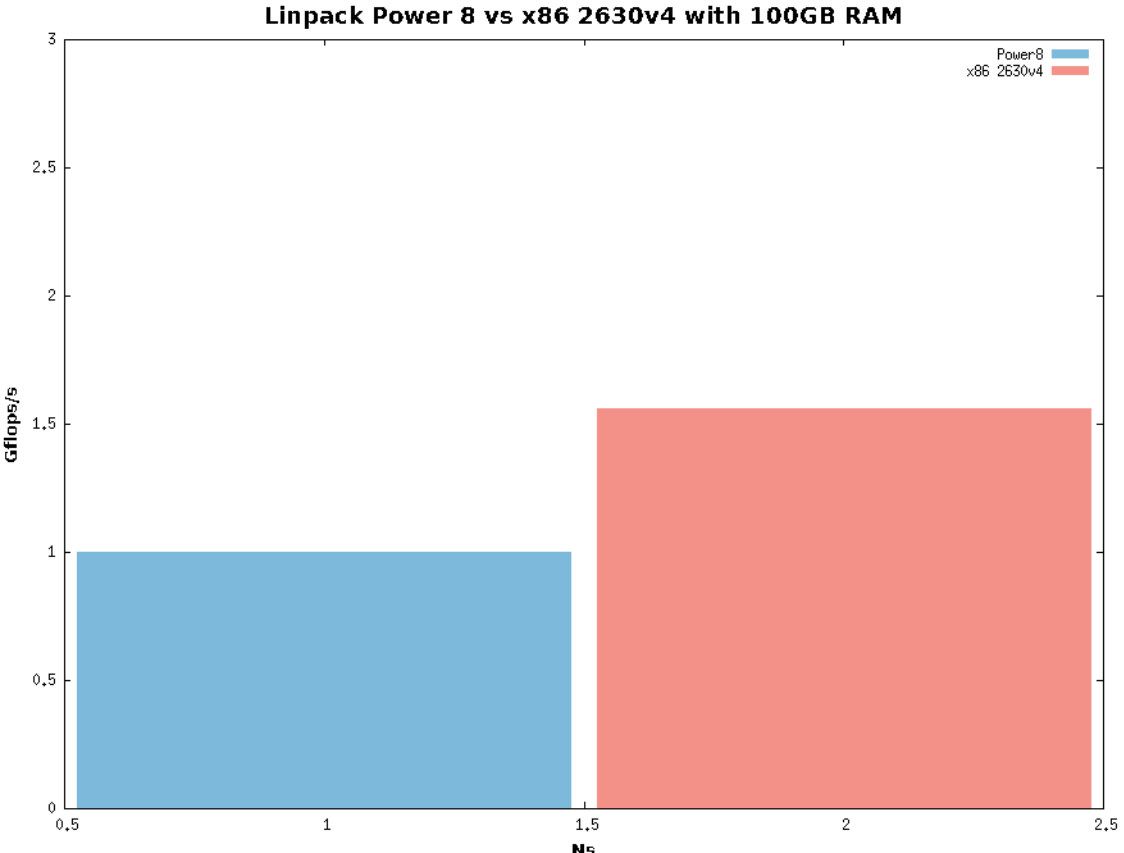

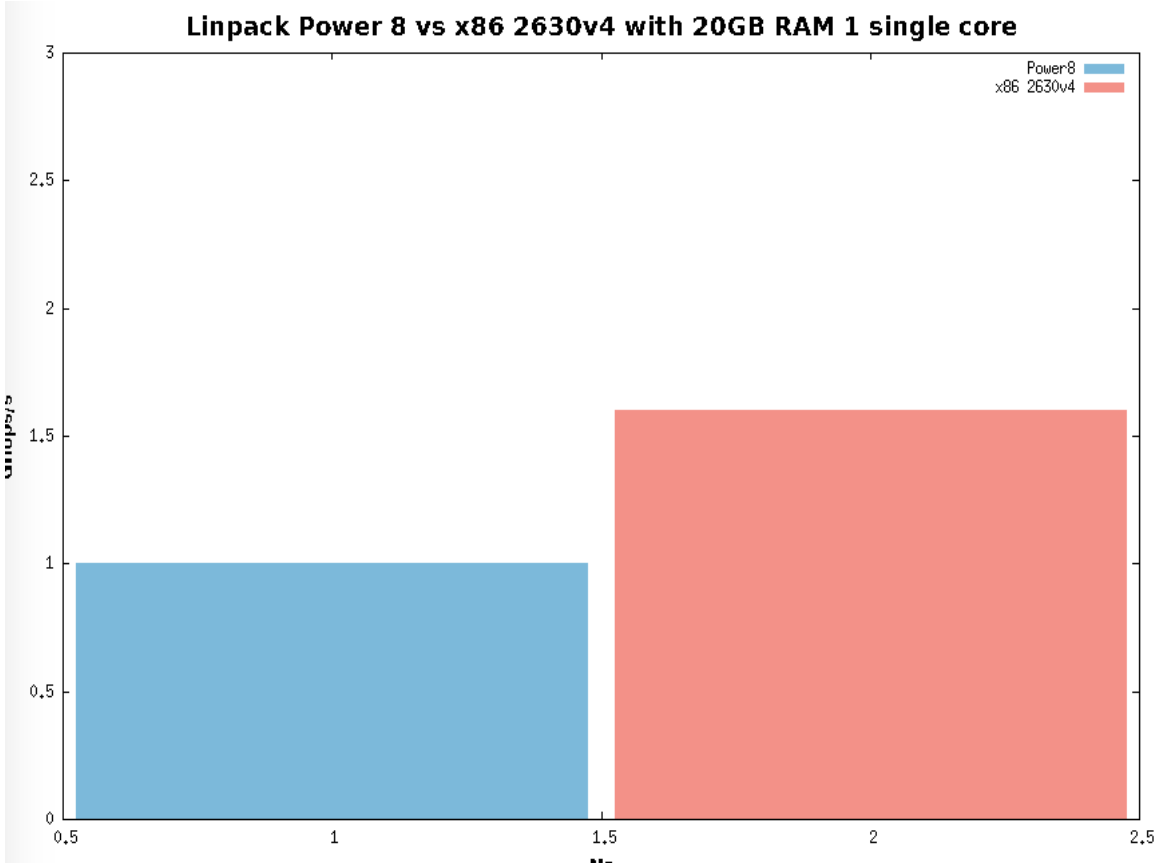

More HPCC TEST Results

http://www.netlib.org/benchmark/hpl/

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

OS: CentOS 7.2.1511 ppc64le Compiler: gcc 4.8.5

Intel(R) Xeon(R) 2CPU 10core per CPU E5-2630v4 @ 2.2GHz, 256 GB RAM

OS: CentOS 7.2.1511 Compiler: Intel Parallel Studio 2017

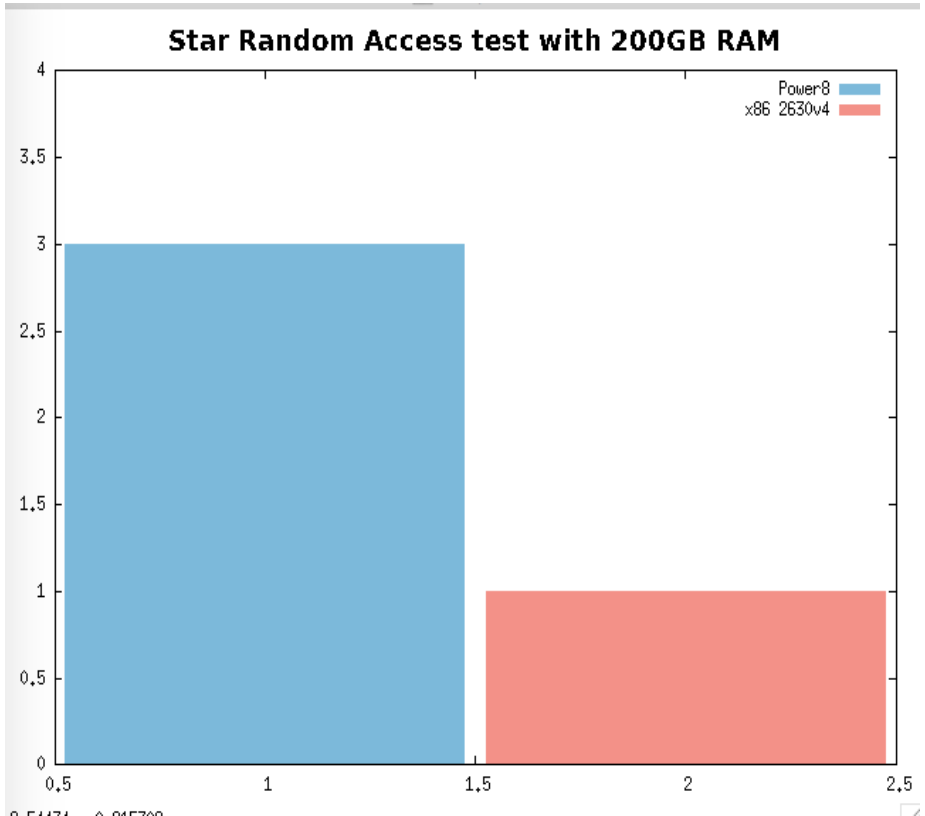

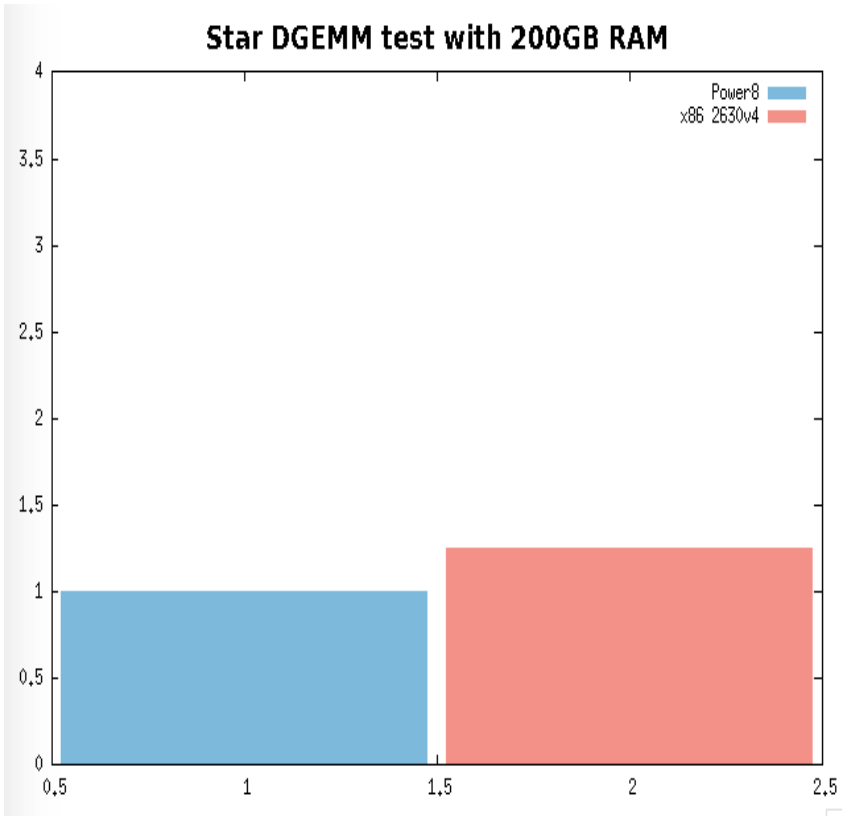

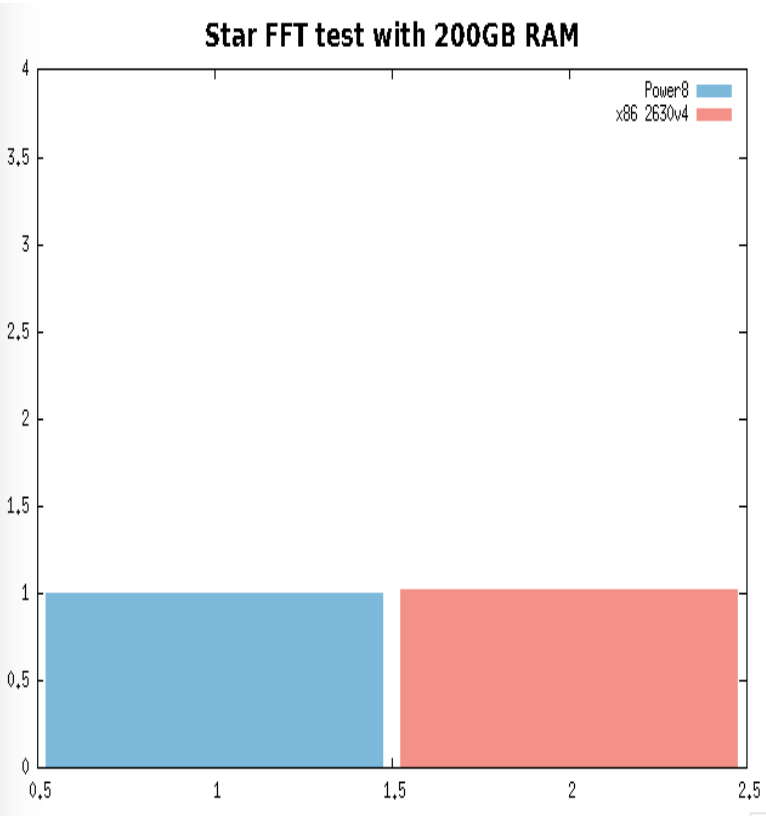

D.A.V.I.D.E. SUPERCOMPUTER

PRACE Awards Third and Final Phase of Pre-Commercial Procurement (PCP)

After successfully completing phase II, during phase III, E4 proposed an innovative design that makes avail of the most advanced technologies, to produce a leading edge HPC cluster showing higher performance, reduced power consumption and ease of use.

http://aec-analisiecalcolo.it/notizie/278-litalia-protagonista-dellinnovazione-grazie-e4-com/

- RESEARCH and PRODUCTION PROJECT
- CUSTOMIZED COMPUTE NODES by WISTRON
- IBM POWER8+ with NVIDIA NVLink
- NVIDIA® TESLA® P100 SXM2
- LIQUID COOLING
- POWER MONITORING & PROFILING, POWER MANAGEMENT, POWER CAPPING & PREDICTION
- CODE PORTING & OPTIMIZATION
- OPEN RACK FORM FACTOR with INTEGRATED PIPING &POWER DISTRIBUTION
- SYSTEM DESIGN, INTEGRATION, TESTING and DELIVERY
- PETAFLOP-CLASS COMPUTING POWER

Scalable and fine-grain power monitoring engine and Job Scheduler Power-Capping for D.A.V.I.D.E.

Beaglebone Black Board

ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

- 1GHz ARM Cortex-A8
- 12bit 8-ch SAR ADC
- PTP HW enabled

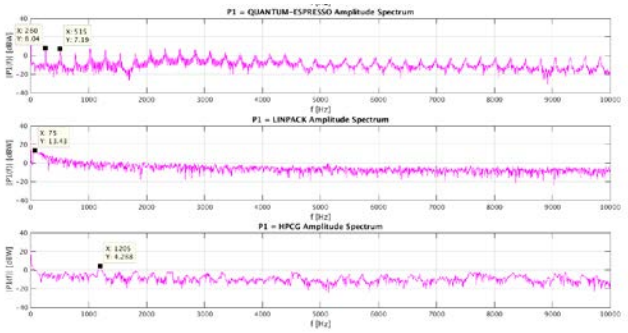

FFT– power spectral analysis

- Without Power Capping

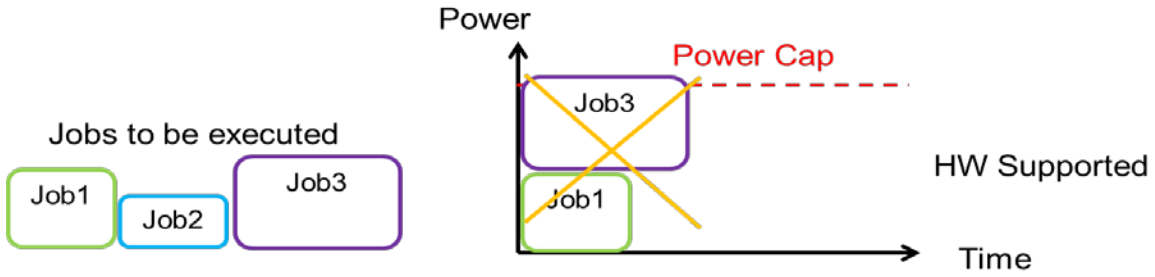

Exceed in power capping

- Our Solution:

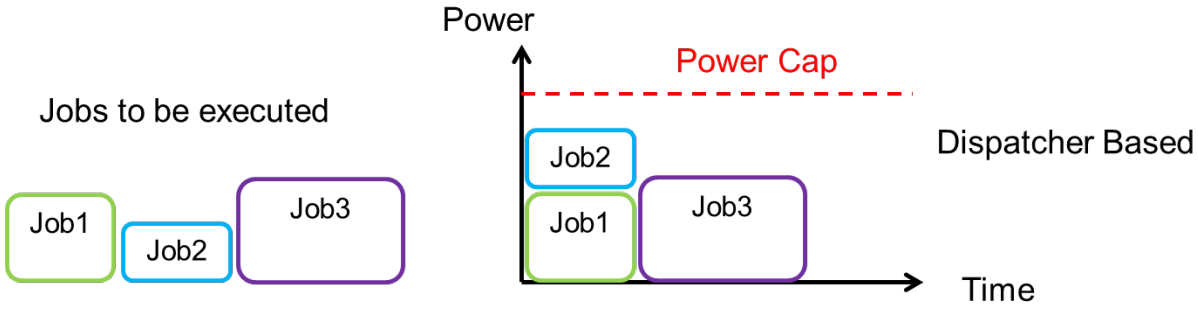

No Job Slowdown

No changes in the accounting

BBB sensors detect thepower consumption and a deep learning system learns the job power consumption, then the scheduler could submit next jobs in base of Power cap limit.

COOPERATION = INNOVATION

MEMBERSHIPS

OpenPOWER Foundation
OpenPower Physic WG
ETP4HPC - HPC European lobby
SPEC
HPC Advisory Council

DESIGN EXASCALE APPLICATIONS

Member of the MAX Center of Excellence of material science (H2020)
Direct financial contribution to port scientific applications on ARM and ARM+GPU architectures

ENERGY AWARENESS COMPUTING

E4 is designing hardware components for high frequency energy sampling
Co-organizer of initiatives, such as COLA workshop, aimed at increasing the knowledge of energy awareness for scientific computing

CONTACTS

Email contacts

info@e4company.com

support@e4company.com

sales@e4company.com

E4 Computer Engineering SpA

Via Martiri della Libertà, 66,

42019 Scandiano (RE) - Italy

Tel. 0039 0522 991811

THANK YOU

E4
COMPUTER
ENGINEERING

BACKUP

Open Power Memory Organizzation

Memory Organization

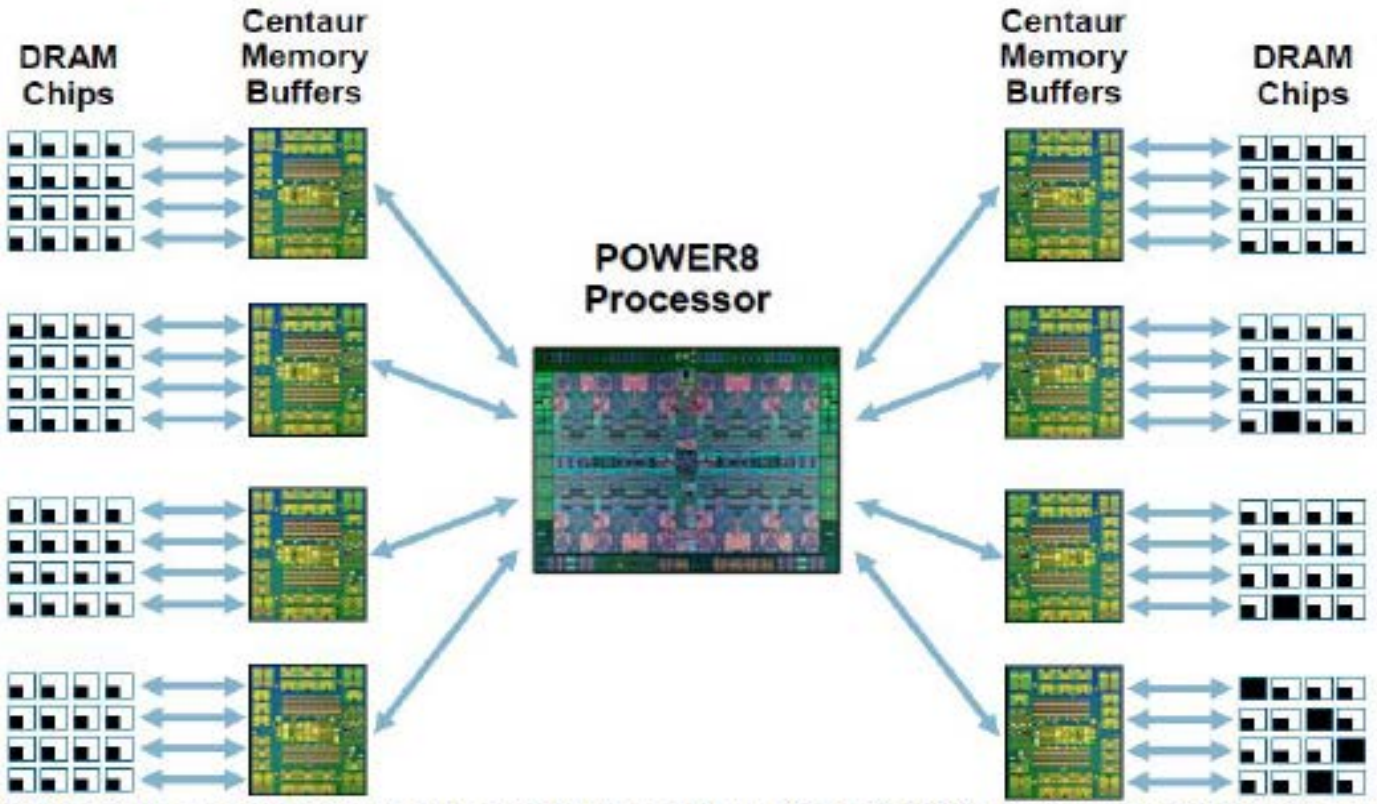

- *Up to 8 high speed channels, each running up to 9.6 Gb/s for up to 230 GB/s sustained*
- *Up to 32 total DDR ports yielding 410 GB/s peak at the DRAM*
- *Up to 1 TB memory capacity per fully configured processor socket*

OUR ECOSYSTEM

END USERS

INFN

CERN

e-geos

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

BAE SYSTEMS

HW INTEGRATORS

DataDirect NETWORKS

EMC²

panasas

sgi

HW MANUFACTURERS

intel

QCT

NVIDIA

applied micro

ARM

CAVIUM

BROCADE

Mellanox

wistron

SOFTWARE VENDORS

ANSYS

Bright Computing

NICE

CD-adapco

PARTNERS

aethia
power computing solutions

ENGINSOFT

Client	NOVARTIS
Industry	Healthcare
Ref.	Riccardo Beltrami – Head of Computational Biology Unit

NOVARTIS

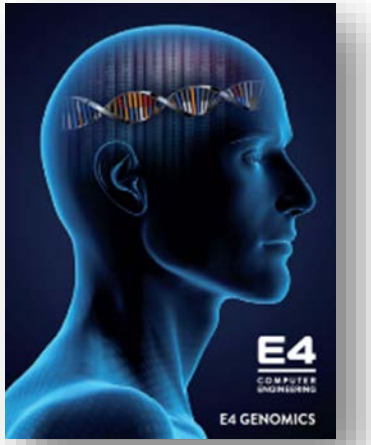

REQUIREMENTS	High Availability HPC Cluster
CHALLENGES	Ensuring faster transfer rate and network interconnection Reducing computing times and obtaining an easy-to-use storage system
SOLUTION	96 dual socket server (1.152 computing cores) Infiniband Switch QDR 40Gbps ports HPC Storage Panasas > 400TB Networking E4 HPC Suite
APPLICATION	HPC Sequencing
KEY FACTORS	Technical skills Post sale support POC Reliability
BENEFITS	Faster processing times More Secure data back-up thanks to HA

Customer BSC PEDRAFORCA CLUSTER

Industry Supercomputing National Centre

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

REQUIREMENTS Custom solution based on low power CPU and GPU accelerators

CHALLENGES Creating an unique prototype with mobile SoC connected to high-end computing GPUs

SOLUTION 78 compute nodes equipped with Tegra 3 SoC, Nvidia K20, Mellanox Infiniband QDR

APPLICATION GPU boosting

KEY FACTORS Low power SoC
Prototyping ability

BENEFITS Accelerated computing at minimum power footprint
First worldwide ARM+GPU prototype
Disruptive innovation

Customer CINECA

Industry Scientific Research

CINECA

REQUIREMENTS

R&D services on
"Whole System Design for Energy Efficient HPC"

CHALLENGES

The goal of this PCP is to procure R&D services that result in highly energy efficient HPC system components that are integrated into an HPC architecture which is capable of providing a floating-point peak performance of up to 1 PFlop/s

APPLICATIONS

Material Science, Climate change, Geophysics, Theoretical Physic

KEY FACTORS

Power Efficiency
1PFlops Pilot system (phase III)

BENEFITS

This PCP aims to foster innovation for economic growth to ensure sustainable high quality public services in Europe and should lead to energy efficient HPC systems suitable for operation within the PRACE infrastructure.

E4 L-U-M-E-A

LUSTRE MADE EASY APPLIANCE

Large storage capacity for fast computing

Easy scalability by providing fast parallel storage

Perfect for large datasets and multiple compute nodes, eliminating downtime and allowing data centers to operate smoothly

Created for HPC users handling large amount of data, Enterprise users managing multiple database, R&D organizations, Engineering depts., weather and climate centers

BENEFITS

- Easy to install & manage
- Linear scalability up to 10/GB throughput
- Cost effective: pay as you grow
- Neutrality by deploying Open source solution
- Capacity: each individual Storage Module can be configured with as much as 1.4PB of raw capacity
- Scale-out the solution with multiple Storage Modules to obtain PB of usable capacity

IBM SPECTRUM SCALE

MANAGE UNSTRUCTURED DATA FOR CLOUD, BIG DATA ANALYTICS, OBJECTS AND MORE

Redefine unified storage to support new-era workloads for file, object, Hadoop and analytics use cases

Achieve new operational efficiency and cost effectiveness—deliver up to 10 times higher performance on the same hardware

Help lower the cost of data retention up to 90 percent through policy-driven automation

Improve application performance by reducing data bottlenecks with flash-based acceleration

Enable collaboration and efficient sharing of resources among global, distributed teams

BENEFITS

- Performance and scalability
- Simplify data management at scale
- Simplified administration
- Global file sharing with active file management
- Synchrounous and asynchrounous disaster recovery
- Multi-protocol support with native access
- Cost-effective information lifecycle management
- End-to-end data reliability, availability and integrity

IBM Spectrum Scale