HPE & E4: Computational Intelligence and Deep Learning for Next-Generation Edge-Enabled Industrial Workflows

Tony De Varco (HPE)
Fabrizio Magugliani (E4 Computer Engineering)

Agenda

WHEN PERFORMANCE MATTERS

- **HPE & E4:** A winning team for the CAE market

  Speakers' introduction
- **E4:** Innovation and effectiveness for CAE workflows
- **HPE:** Solutions for CAE workloads
- **HPE & E4:** High tech and customer care

HPE & E4: a winning team for the CAE market

WHEN PERFORMANCE MATTERS

Tony De Varco

Tony DeVarco is the HPC, Manufacturing Segment Manager at Hewlett Packard Enterprise. Tony is responsible for strategy and market analysis, competitive analysis, definition of value proposition, partner ecosystems, segment trends and knowledge in the Virtual Product Design (VPD) and Manufacturing segment including Computer Aided Engineering (CAE) and Electronic Design Automation (EDA) workflows.

anthony.devarco@hpe.com

Fabrizio Magugliani

Fabrizio Magugliani is Horizon Europe and EuroHPC Project Manager at E4 Computer Engineering. Fabrizio coordinates large-scale projects, including the development of complex HW equipment and the integration of the SW. An Aerospace Engineer as background, Fabrizio is representative of E4 in the OpenFOAM HPC Technical Committee, in EPI and in ETP4HPC.

fabrizio.magugliani@e4company.com

Innovation and Effectiveness for CAE Workflows

CAE Conference 2021

Mixing the right components

WHEN PERFORMANCE MATTERS

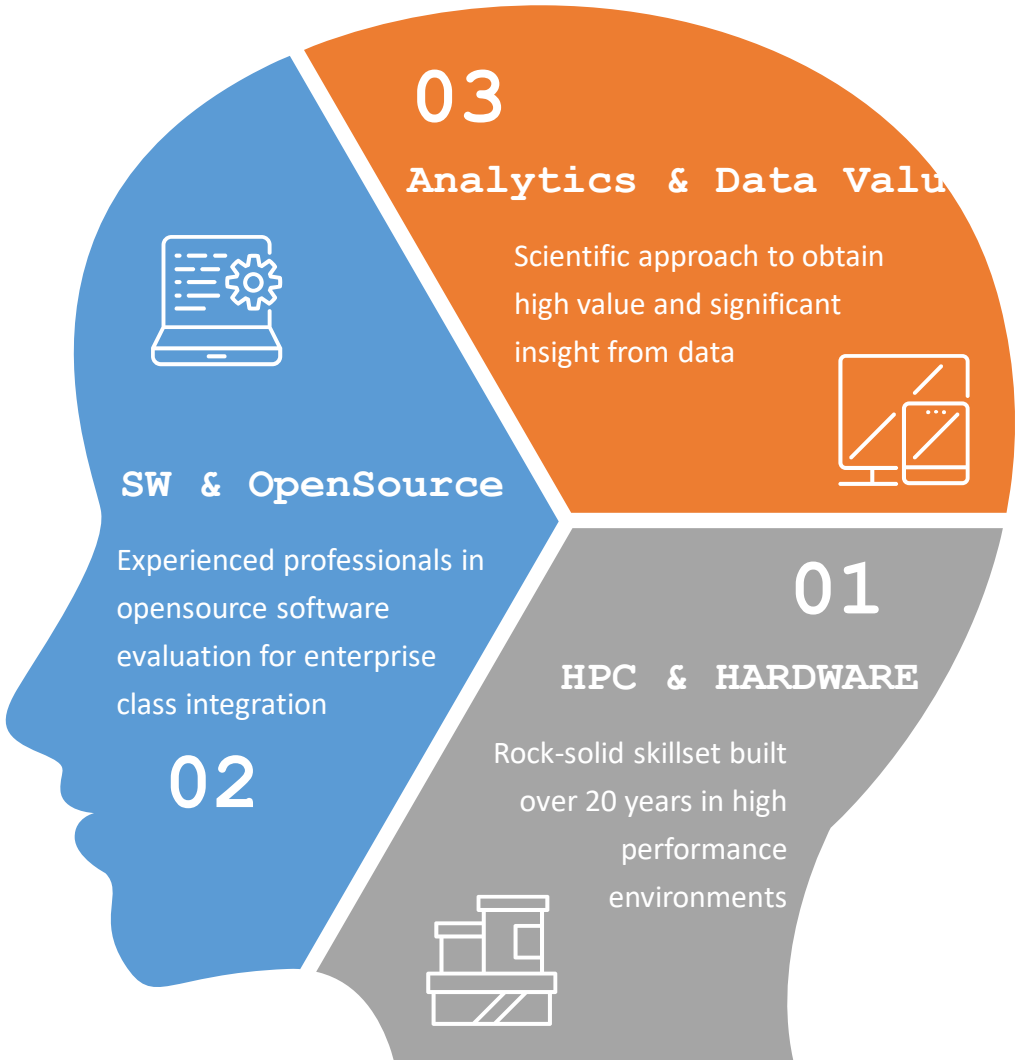

E4 History

WHEN PERFORMANCE MATTERS

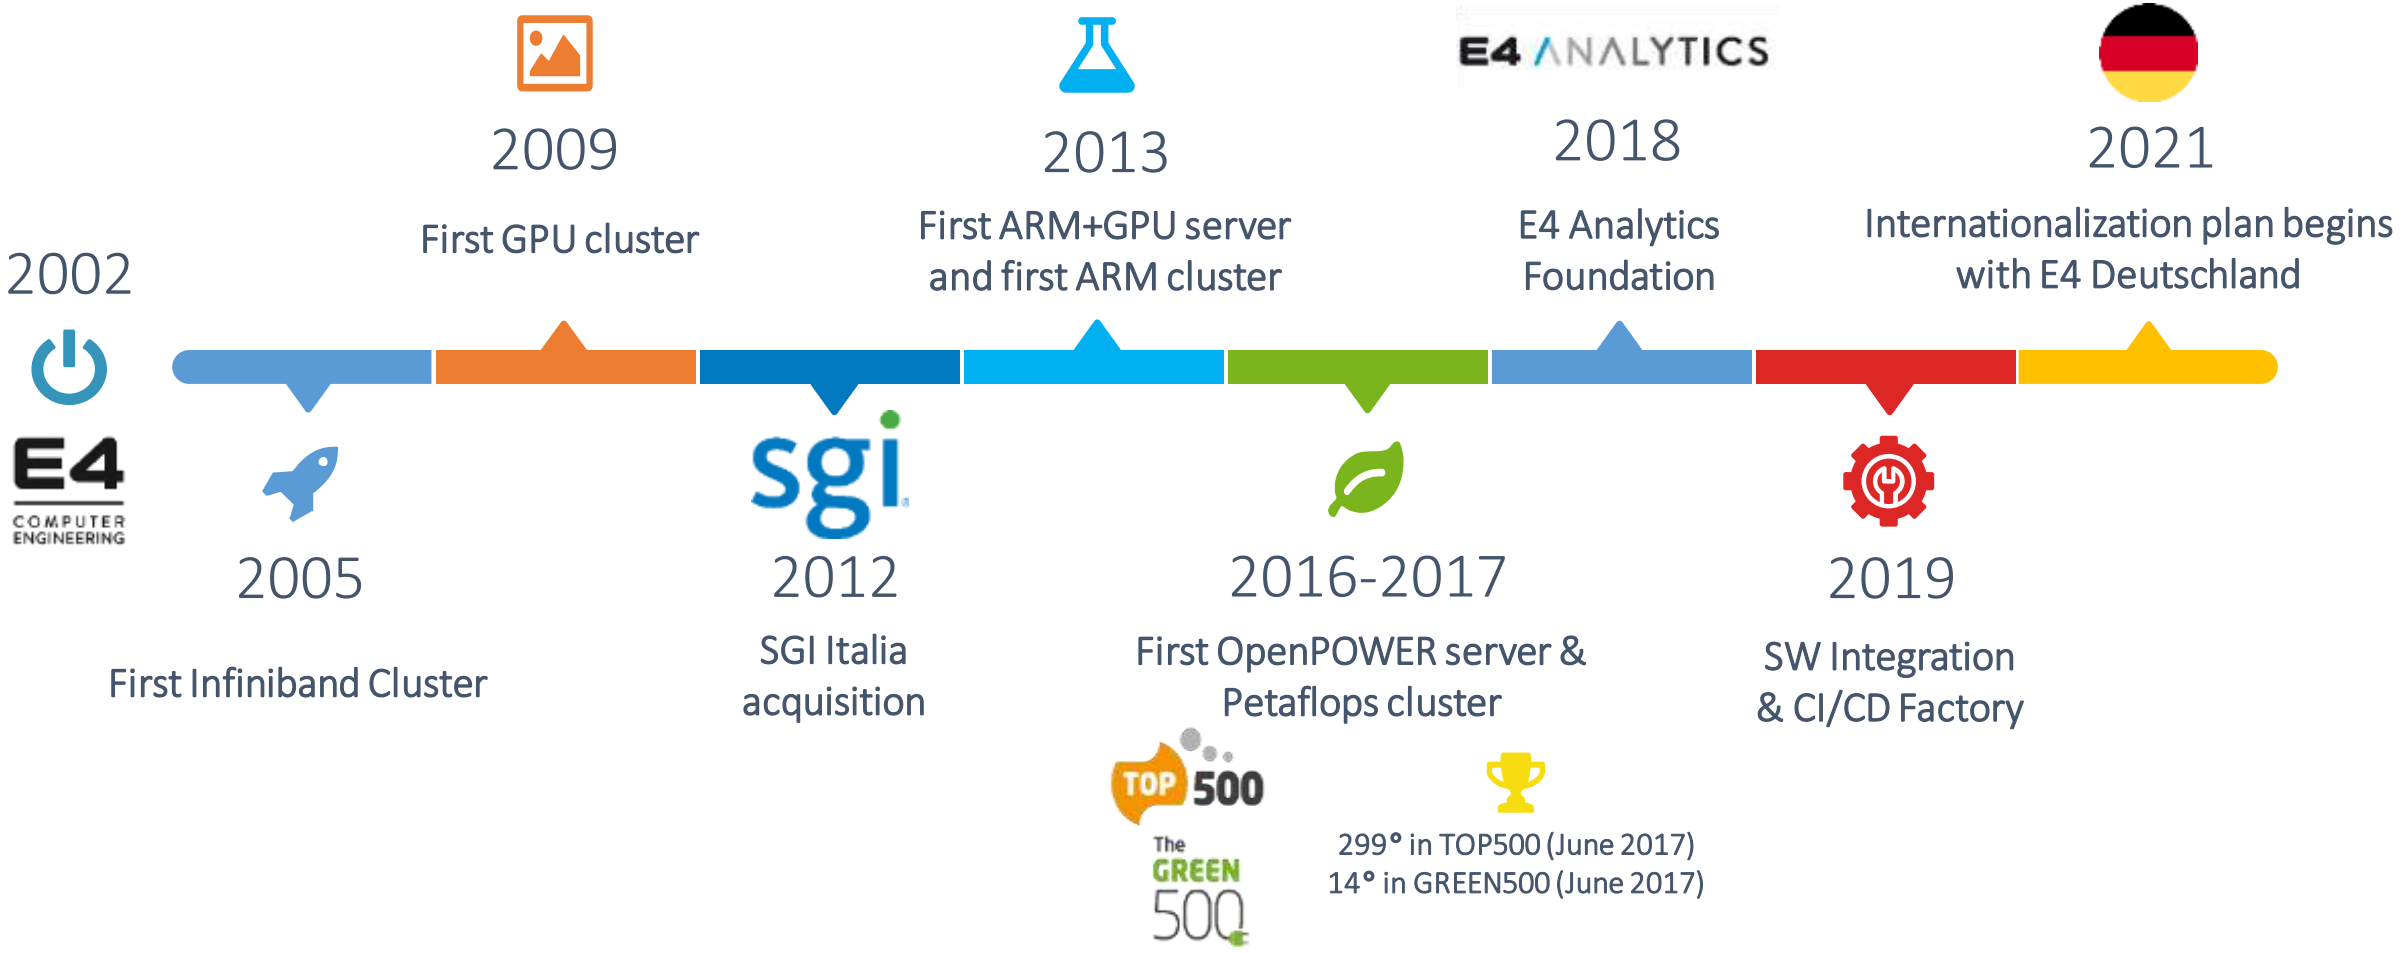

E4 in numbers

WHEN PERFORMANCE MATTERS

290x

Acceleration factor in healthcare image processing using GAIA

>40%

YoY revenues growth from 2019

+10.000

Systems units provided to CERN Geneve

+400

Active customers (last 3 yr)

>50%

Tech – R&D People

>25Gbps

Data throughput achieved with USTI accelerated storage saturating the network

>20%

Operational yearly costs saving with automation provided by Fluctus

+10

Data Scientist ready-to-use frameworks integrated in E4 Analytics Studio

E4 Growth

WHEN PERFORMANCE MATTERS

Revenues

People

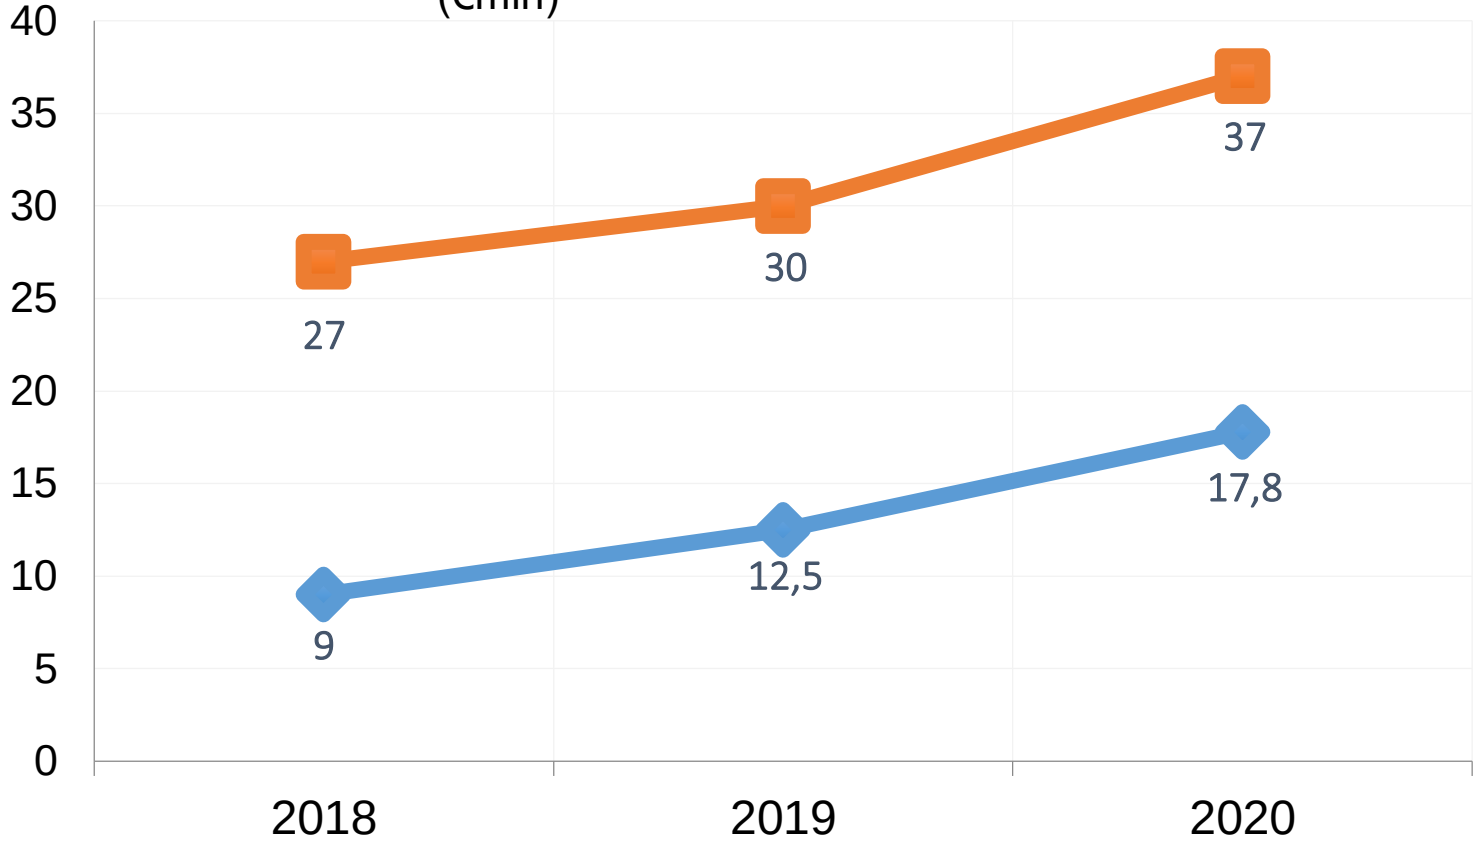

E4 Computer Engineering/memberships

WHEN PERFORMANCE MATTERS

Member of the Steering Board
http://www.etp4hpc.eu

ANDREAS: Artificial intelligence traiNing scheDuler foR disaggrEgAted resource clusterS
Value Chain Oriented and Interdisciplinary Technology Transfer EXperiments (TTX)-
https://www.tetramax.eu/ttx/funded-projects/#/

Member of the Consortium
http://european-processor-initiative.com

Partnership Programme
https://www.openfoam.com/

Member of CERN openlab
https://openlab.cern/

Member of the OEHI (Open Edge and HPC Initiative)
http://www.open-edge-hpc-initiative.org/

Member of the MaX Center of Excellence
http://www.max-centre.eu/

Member of HiPEAC
https://www.hipeac.net/

THE EUROPEAN PROCESSOR INITIATIVE (EPI/SGA1)

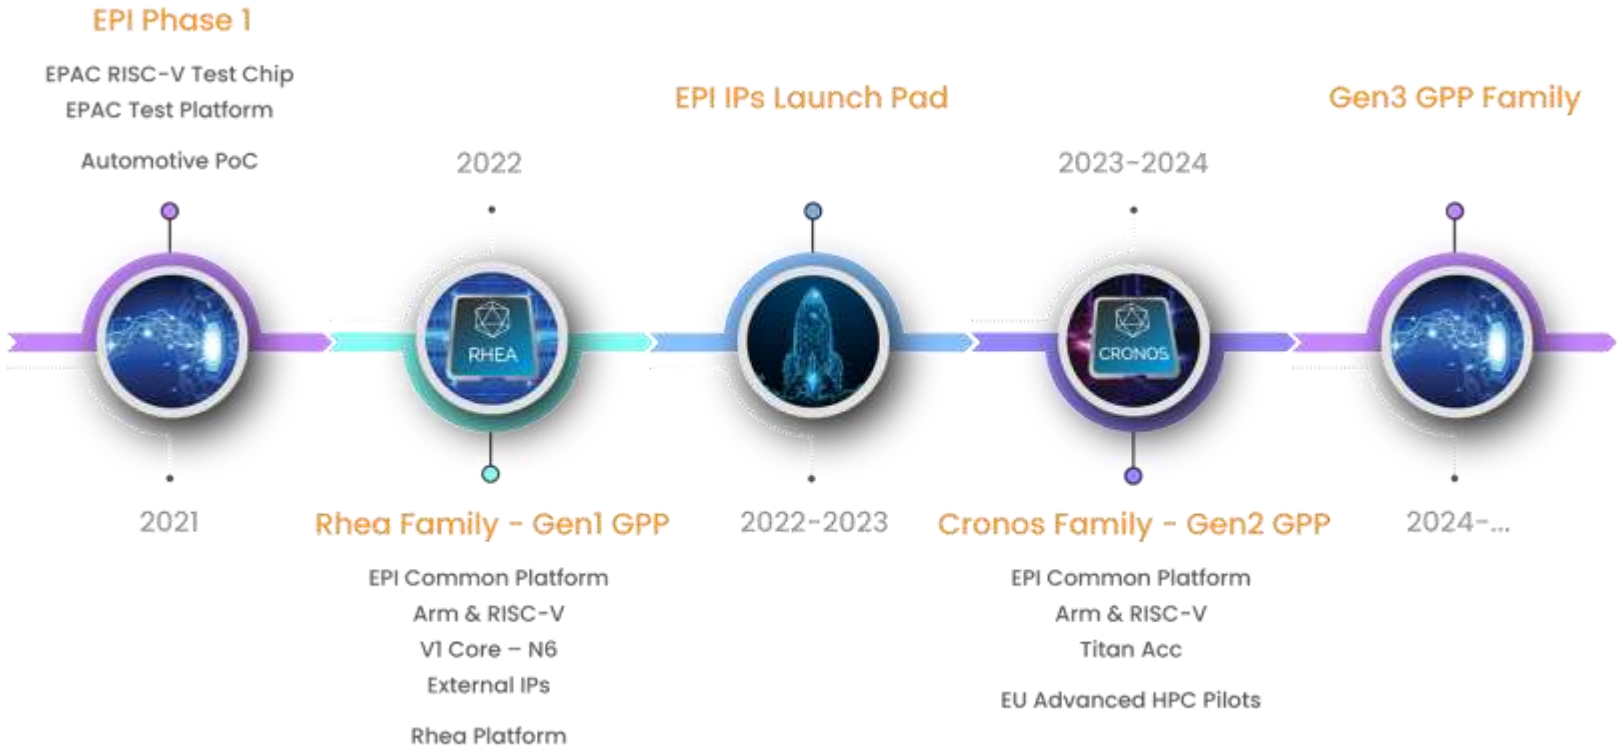

Implemented under the first stage of the Framework Partnership Agreement signed by the Consortium with the European Commission (FPA: 800928)

EPI/SGA1 IN THE PRESS

Press release: EPI EPAC1.0 RISC-V core boots Linux on FPGA
March 9, 2021
https://www.european-processor-initiative.eu/epi-epac1-0-risc-v-core-boots-linux-on-fpga/

Press release: EPI EPAC1.0 RISC-V Test Chip Taped-out
June 1st, 2021
https://www.european-processor-initiative.eu/epi-epac1-0-risc-v-test-chip-taped-out/

AS OF EARLY SEPTEMBER…

AND AS OF LAST WEEK…

AND AS OF "NOW"!

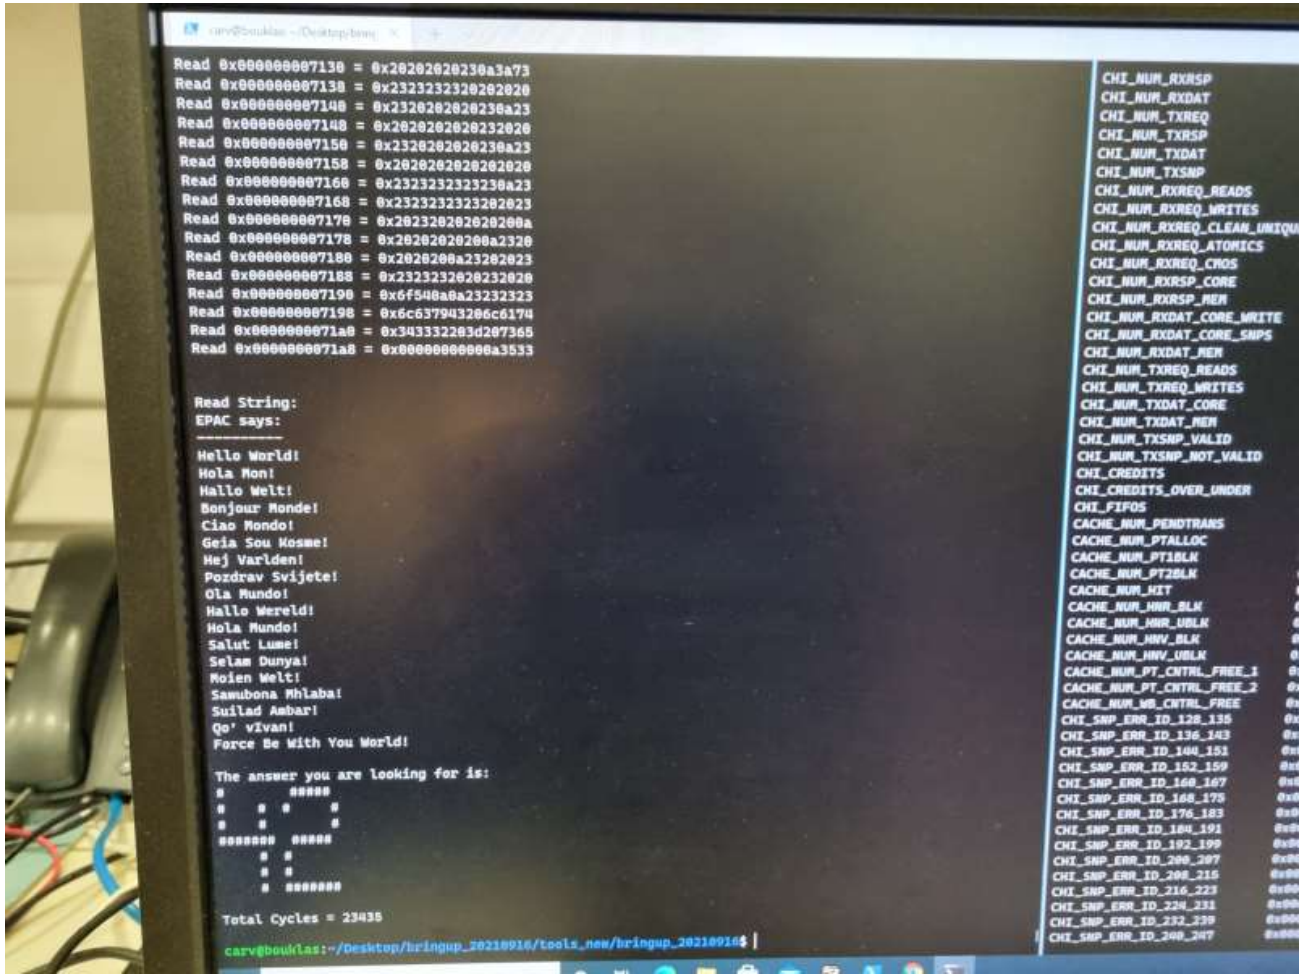
Read 0x0000000007130 = 0x20202020230a3a73
Read 0x0000000007138 = 0x2323232320202020
Read 0x0000000007140 = 0x2320202020230a23
Read 0x0000000007148 = 0x2020202020232020
Read 0x0000000007150 = 0x2320202020230a23
Read 0x0000000007158 = 0x2020202020202020
Read 0x0000000007160 = 0x2323232323230a23
Read 0x0000000007168 = 0x2323232323202023
Read 0x0000000007170 = 0x202320202020200a
Read 0x0000000007178 = 0x20202020200a2320
Read 0x0000000007180 = 0x2020200a23202023
Read 0x0000000007188 = 0x2323232020232020
Read 0x0000000007190 = 0x6f540a0a23232323
Read 0x0000000007198 = 0x6c637943206c6174
Read 0x00000000071a0 = 0x343332203d207365
Read 0x00000000071a8 = 0x0000000000a3533
Read String:
EPAC says:
Hello World!
Hola Mon!
Hallo Welt!
Bonjour Monde!
Ciao Mondo!
Geia Sou Kosme!
Hej Varlden!
Pozdrav Svijete!
Ola Mundo!
Hallo Wereld!
Hola Mundo!
Salut Lume!
Selam Dunya!
Moien Welt!
Sawubona Mhlaba!
Suilad Ambar!
Qo' vIvan!
Force Be With You World!
The answer you are looking for is:
Total Cycles = 23435
carv@bouklas:~/Desktop/bringup_20210916/tools_new/bringup_20210916$
CHI_NUM_RXRSP
CHI_NUM_RXDAT
CHI_NUM_TXREQ
CHI_NUM_TXRSP
CHI_NUM_TXDAT
CHI_NUM_TXSNP
CHI_NUM_RXREQ_READS
CHI_NUM_RXREQ_WRITES
CHI_NUM_RXREQ_ATOMICS
CHI_NUM_RXREQ_CMOS
CHI_NUM_RXRSP_CORE
CHI_NUM_RXRSP_MEM
CHI_NUM_RXDAT_CORE_WRITE
CHI_NUM_RXDAT_CORE_SNPS
CHI_NUM_RXDAT_MEM
CHI_NUM_TXREQ_READS
CHI_NUM_TXREQ_WRITES
CHI_NUM_TXDAT_CORE
CHI_NUM_TXDAT_MEM
CHI_NUM_TXSNP_VALID
CHI_NUM_TXSNP_NOT_VALID
CHI_CREDITS
CHI_CREDITS_OVER_UNDER
CHI_FIFOS
CACHE_NUM_PENDTRANS
CACHE_NUM_PTALLOC
CACHE_NUM_PT1BLK
CACHE_NUM_PT2BLK
CACHE_NUM_HIT
CACHE_NUM_HNR_BLK
CACHE_NUM_HNR_UBLK
CACHE_NUM_HNV_BLK
CACHE_NUM_HNV_UBLK
CACHE_NUM_PT_CNTRL_FREE_1
CACHE_NUM_PT_CNTRL_FREE_2
CACHE_NUM_WB_CNTRL_FREE

E4 Markets

WHEN PERFORMANCE MATTERS

CAE & AI/DL/ML: integrating intelligence in the w

WHEN PERFORMANCE MATTERS

When engineers discuss AI, they're usually focusing on AI *models,* but AI is much more. The concept of "AI/DL/ML" spans four steps within a workflow: data preparation, modelling, simulation and testing, and deployment.

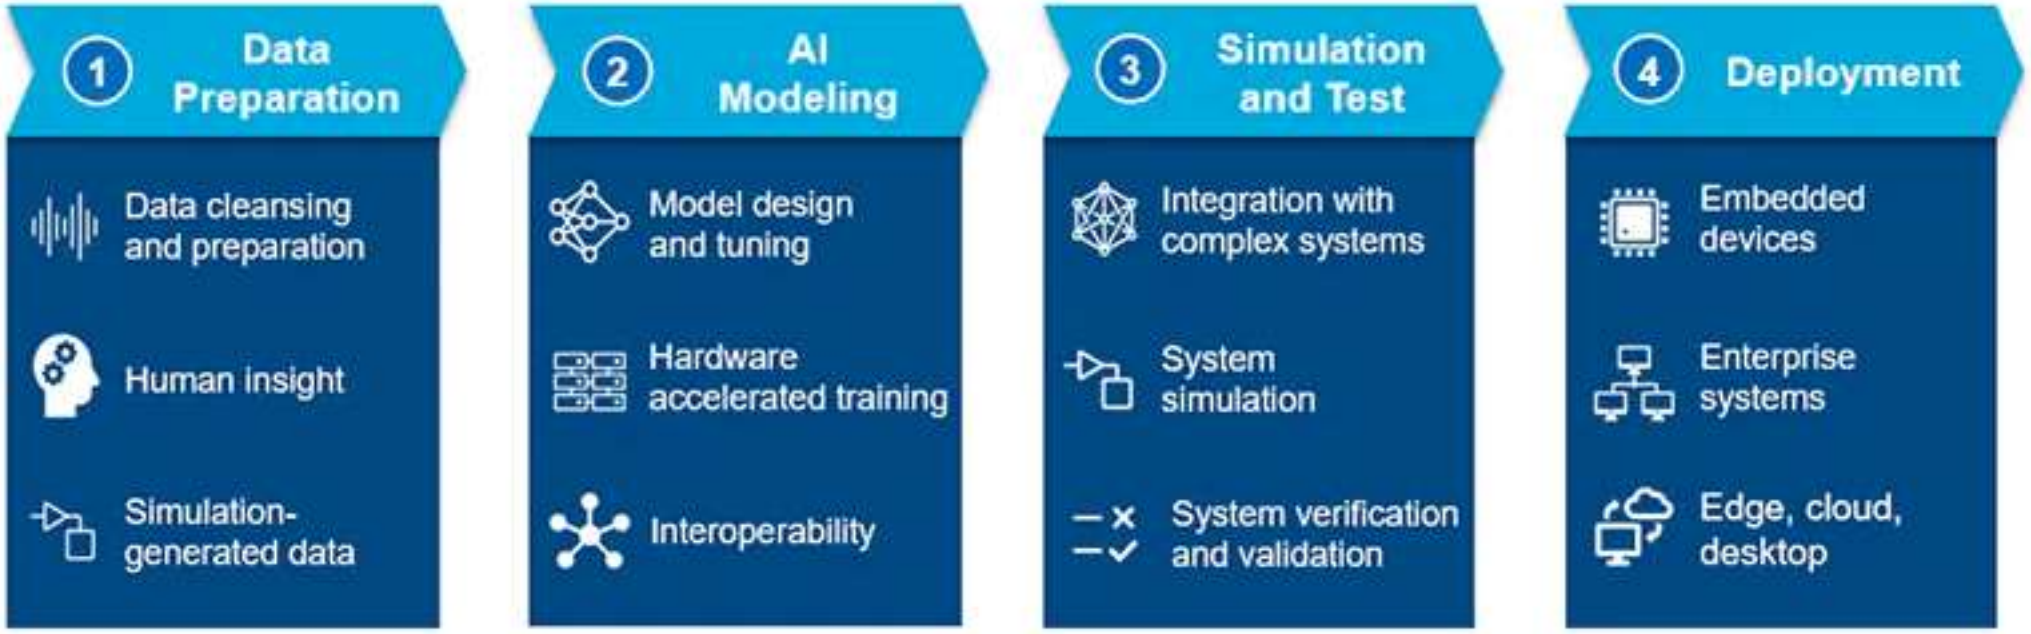

Key takeaway: Engineering firms are integrating AI/DL/ML into their projects, both to improve their results and to remain ahead of their competitors

CAE & IoT: integrating data in the design worl

WHEN PERFORMANCE MATTERS

- The Internet of Things (IoT) is currently transforming the global economy and promises to continue powering its evolution.
- IoT-enabled devices are constantly collecting performance data. For example, production machinery provides data that can be used to calculate overall equipment effectiveness (OEE) on the factory floor. Or wearable devices might provide easy-to-understand health metrics for consumers.
- On the production side, all this data shows how products actually perform in the real world. Engineers can use the data to proactively improve the design of future products. Using data in this way can result in shorter, less expensive product development cycles by enabling the creation of better prototypes from the very start of the design process.
- The constant stream of data offers a treasure trove for engineers who wish to make data-driven product improvements using real-time, real-environment information, it also presents a challenge in terms of data storage and analysis.

Key takeaway: In the age of IoT, simulations (on-prem or cloud-based) are more important than ever.

Industry 4.0: from the idea to the final produ

WHEN PERFORMANCE MATTERS

- Engineers can no longer afford the ‘build it and tweak it’ approach that has long characterized many design projects as we move towards zero prototyping
- Manufacturers today need to implement rigorous systems-design processes that accommodate the complexities of developing multi-disciplinary systems, with high-fidelity virtual prototypes, or ‘digital twins’, at the core of their development process
- HPC infrastructures, optimized to deliver fast response time within a seamless continuity of functions, enable engineers, designers and enterprise to reduce development time, improve products’ quality and features and maximize the ROI
- On-site scalable infrastructures, complemented by high-scalability applications such as ANSYS, represent the ideal infrastructure for implementing and reaping the full-scale benefits of Industry 4.0

Key takeaway: The convergence of mechatronic, cyber-physical technologies, advances in data management, artificial intelligence (AI), machine-learning (ML), and communications via the Internet of Things shapes today’s industrial manufacturing processes.

The enabling solutions

WHEN PERFORMANCE MATTERS

E4 COMPUTER ENGINEERING

E4/HPE CAE Conference 2021

CAE solution

- Scalable cluster
- High frequency CPU
- Large core count
- Large memory capacity
- Fast memory bandwidth
- Greater I/O performance
- Lower network latency
- Higher network bandwidth
- Better RAS

HPE Apollo 2000 Gen10 Plus System for AMD and Intel CPU's

- Up to 4 x AMD EPYC CPU-based servers per 2U chassis
- 2 x AMD EPYC 7002 or 7003 series processors per server
- Up to 64 cores and 128 threads per CPU
- 2 TB memory per server (8 TB in 2U) - 16 x 128 GB
- Eight memory channels for superior throughput
- Up to 3200 MT/s DDR4 memory
- Up to 32 MB shared L3 cache per core (7003 series)
- Up to 16 MB shared L3 cache per core (7002 series)
- Hot Plug SFF SATA/SA and NVMe storage options
- Comprehensive management tools (APM/RCM)
- Security anchored in HPE iLO 5 and Silicon Root of Trust
- 2 x 3000W power supplies, N+N redundancy
- Enhanced thermal efficiency for HPC workloads
- Optional internal RAID controllers

HPE Apollo 2000 Gen10 Plus System front view with multiple storage options

HPE Apollo 2000 Gen10 Plus System rear view with up to 4 hot-pluggable dual-processor servers per chassis for maximum density and flexibility

E4 is hardware- and vendor-agnostic selecting the best technology based on the customer'

Hewlett Packard Enterprise

HPE SOLUTIONS FOR CAE WORKLOADS

MANUFACTURER'S PRODUCT LIFE CYCLE

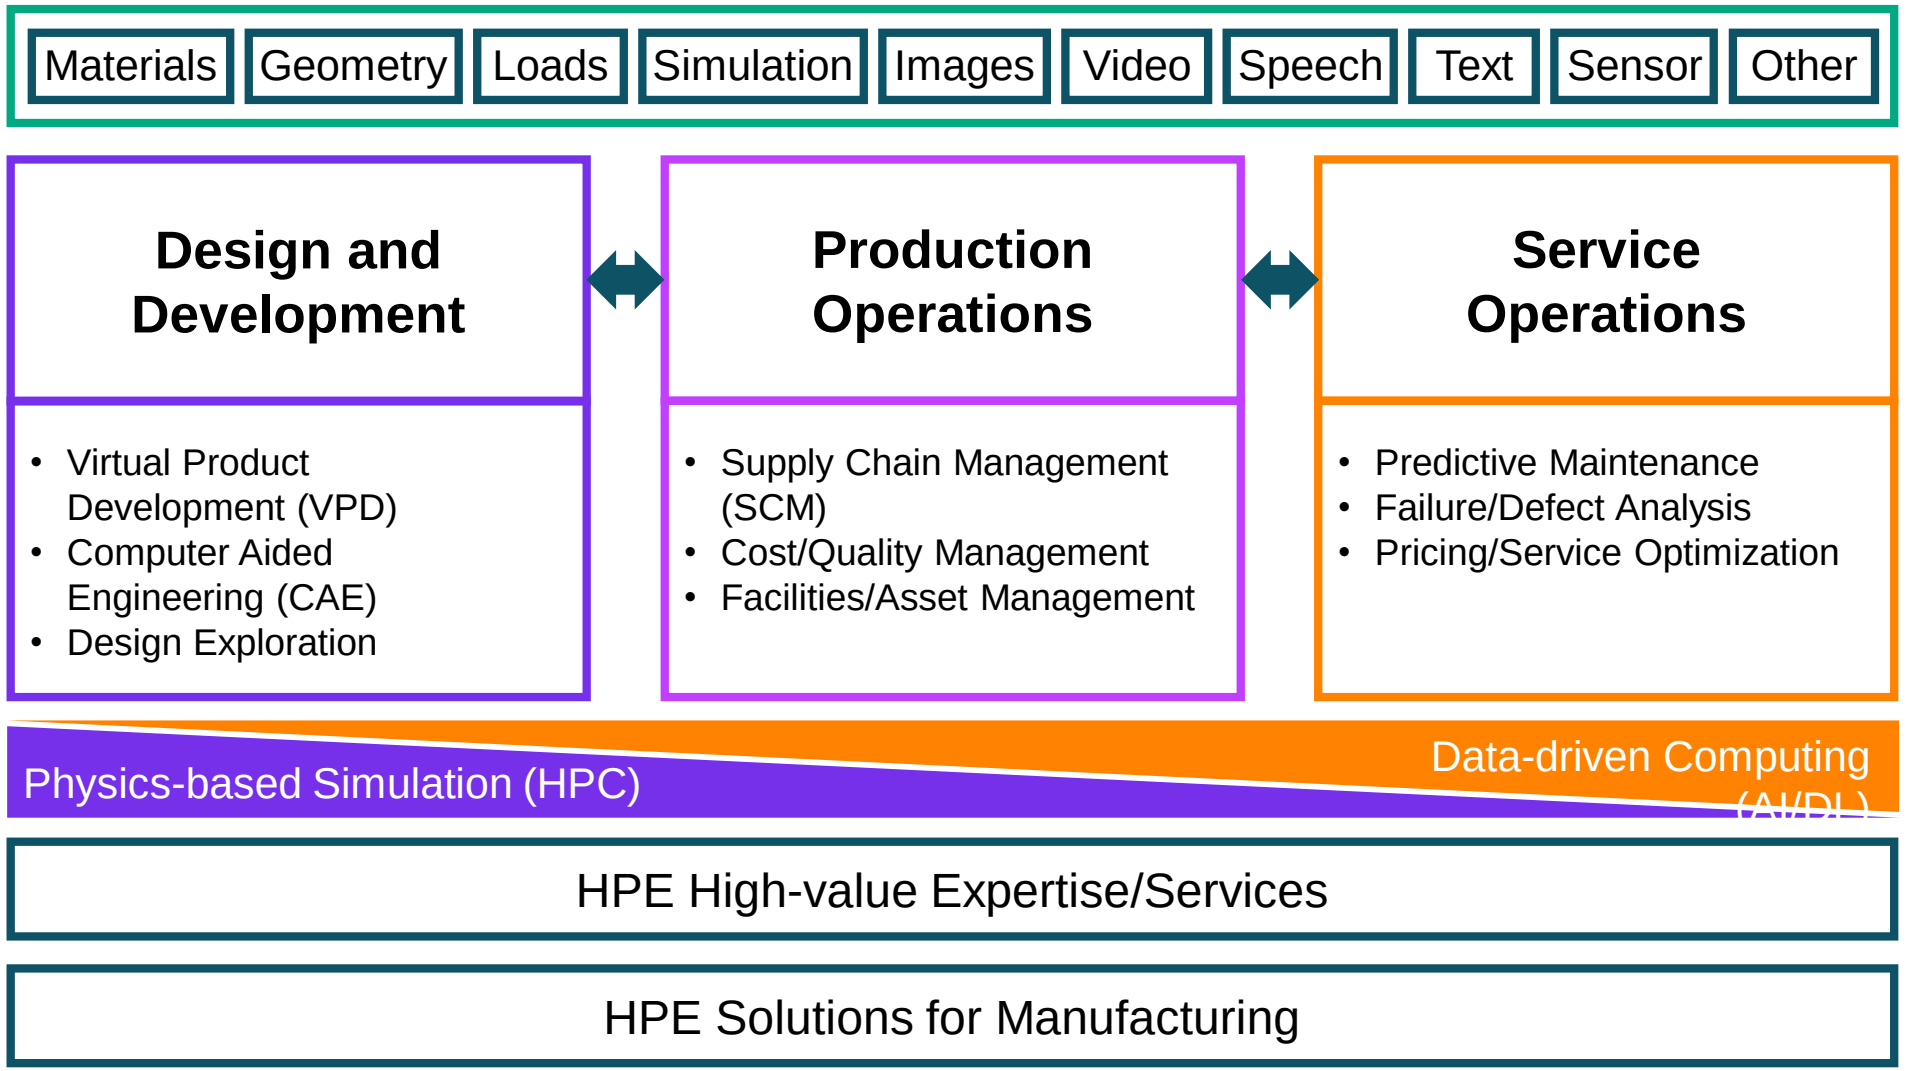

MANUFACTURING AND "DIGITAL TWIN" WORKFLOW

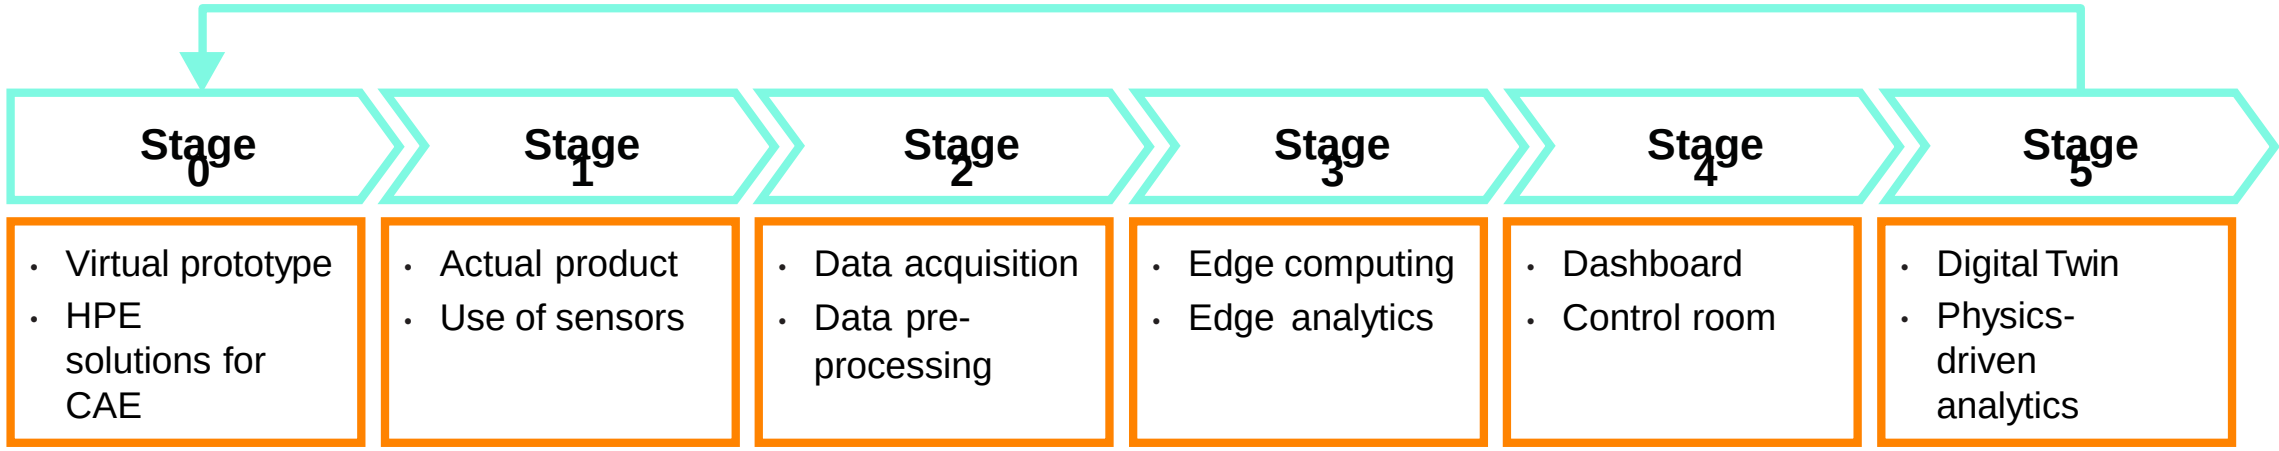

REAL TIME MONITORING AND ROOT CAUSE ANALYSIS AND DIGITAL TWIN

Preventive maintenance avoids costly downtime

Stage 0 | Stage 1 | Stage 2 | Stage 3 | Stage 4 | Stage 5

Ansys

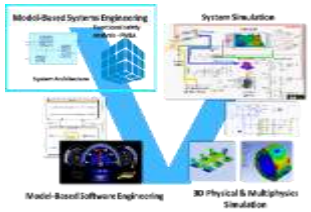

Key Failure Modes/Patterns
Sensor Placement
Design for Operations

Hewlett Packard Enterprise

HPC Solutions for
CAE Workloads

FLOWSERVE

Real Time Augmented
Reality Application

NATIONAL INSTRUMENTS

NI cRIO

Data Acquisition &
Pre-processing

The Edge

aruba

iWap

Hewlett Packard Enterprise

Edgeline EL 1000
for Analytics
(ThingWorx)

Cloud or
DC
ThingWorx

Dashboard
(Control Room)

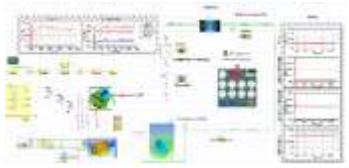

Simulated screen mage courtesy of Ansys

Deep Insights
Physics Driven Analytics
Optimize/Predict with Confidence
Digital Twin

Root Cause Analysis, Corrective Action Process, and augmented reality for service support

ptc

Ansys FLOWSERVE

CAE DISCIPLINES

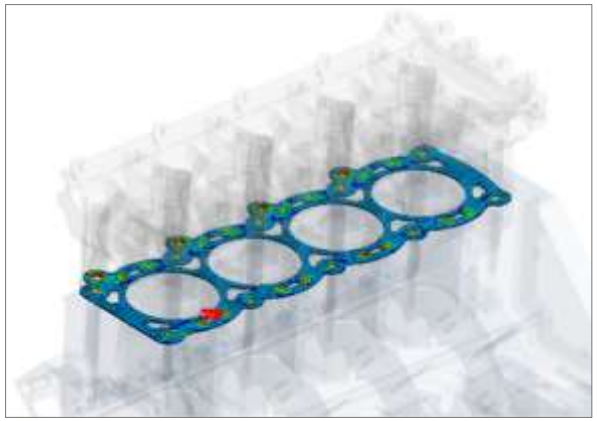

Computational structural mechanics (CSM) for implicit FEA

Simulate the strength and vibration characteristics of product.

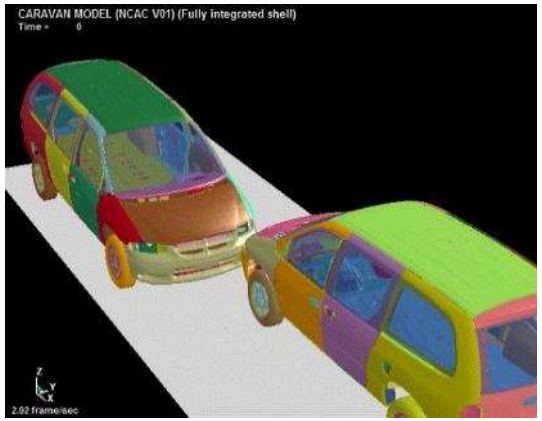

Computational structural mechanics (CSM) for explicit FEA

Simulate the shock impact of products over a short duration.

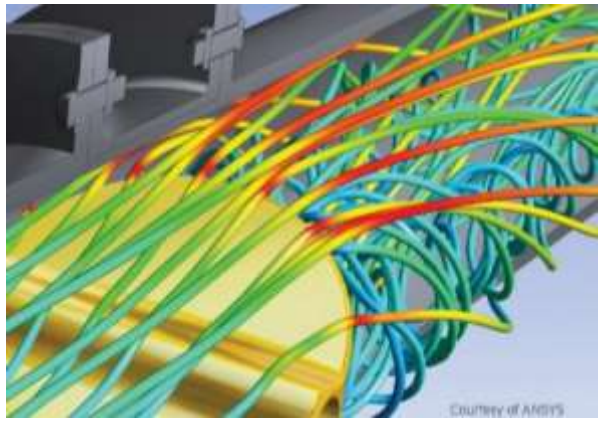

Computational fluid dynamics (CFD)

Simulate aerodynamics; cooling; and mixing of fluids such as air, water, and chemicals.

Computational electromagnetics (CEM)

Simulate radar signature/scattering to assess and prevent detection and identification, antenna performance, ASIC package simulations.

Images courtesy of Ansys

MEETING MANUFACTURING CHALLENGES WITH CAE

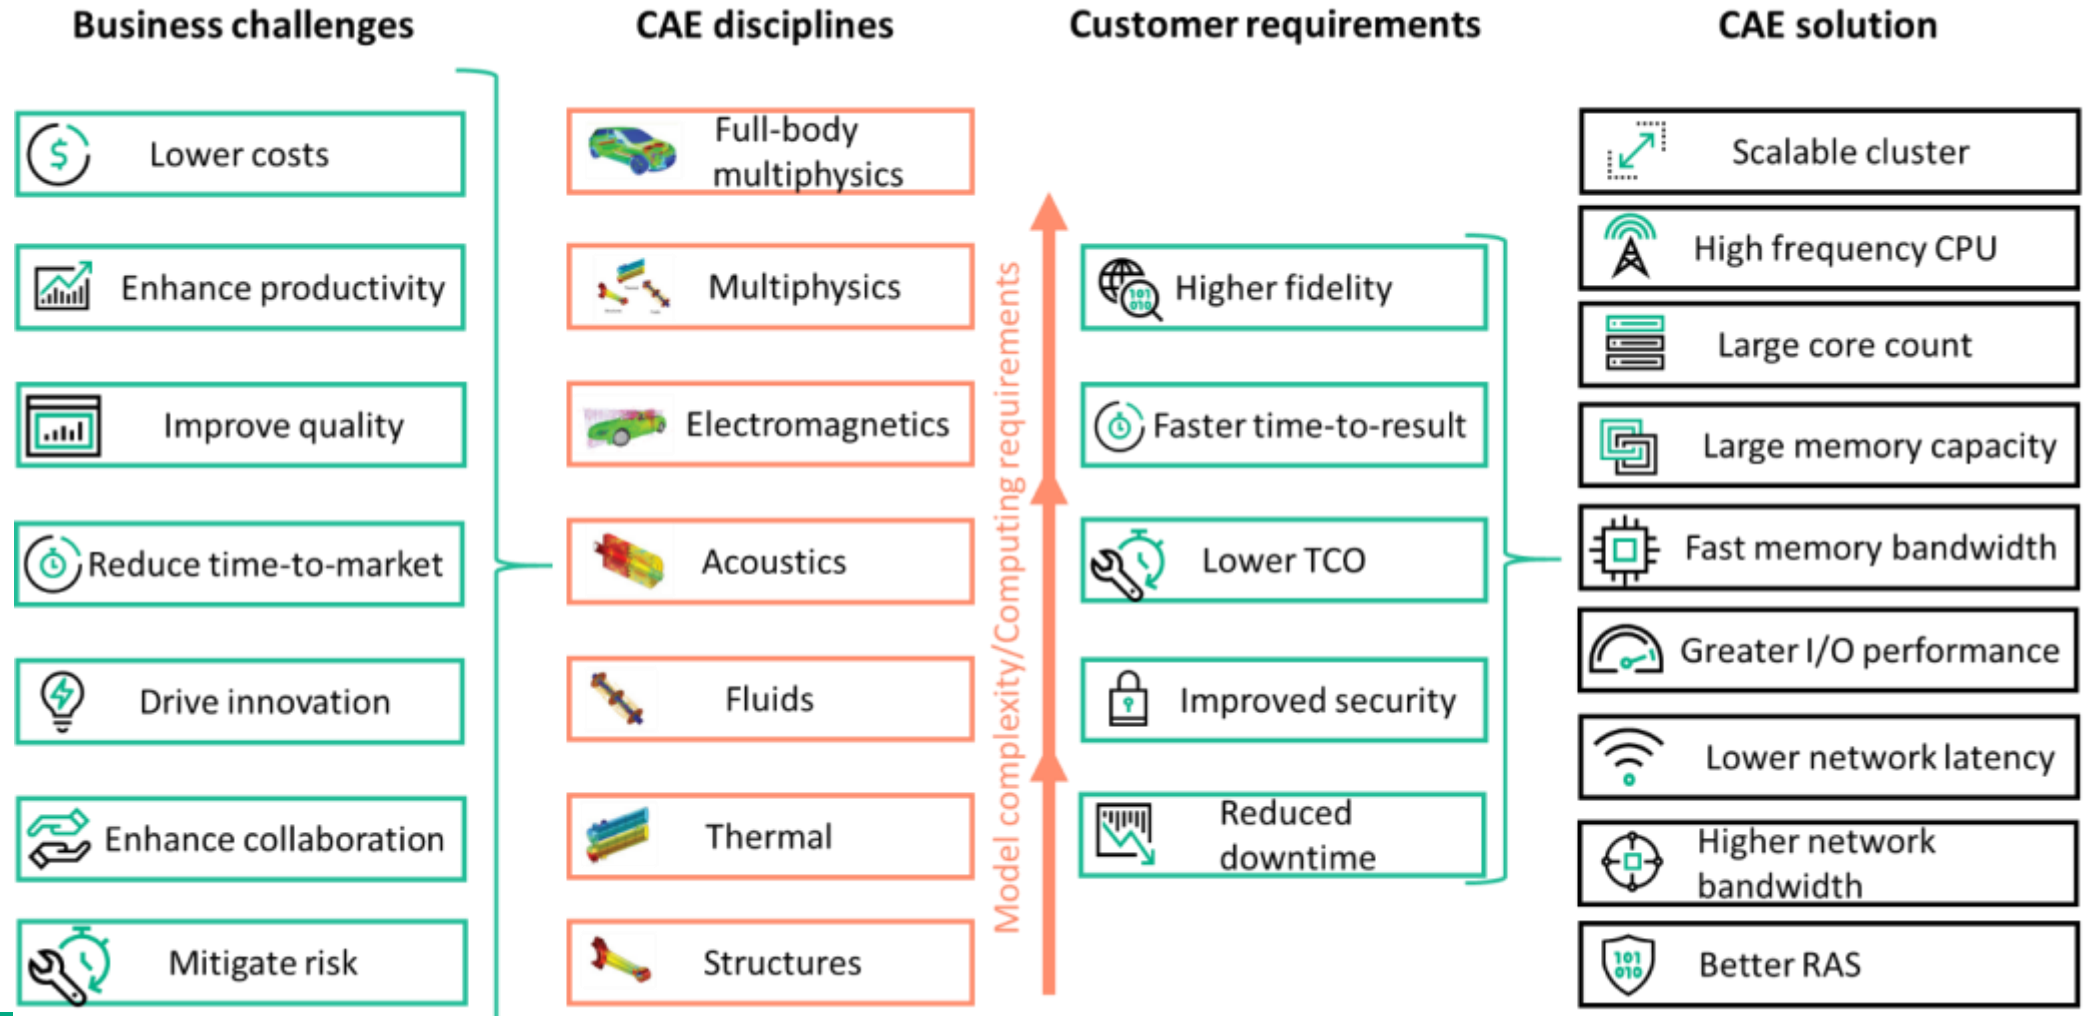

HPE'S ENABLING TECHNOLOGIES/BUILDING BLOCKS

Enterprise HPC, Deep Learning, Inference and Virtualization

HPE Apollo 2000 Gen10 Plus System for AMD and Intel CPU's

System performance and optimization
- 2x compute density of traditional 1U server
- Expanded power capabilities & plug and play Direct Liquid Cooling options
- Software development and application acceleration tools for application performance at scale

Flexible Scale-out Building Blocks
- Storage and I/O flexibility
- Right size building blocks with future proof scalability
- Comprehensive software portfolio to accommodate any workload

Comprehensive server security and management
- Secure from the start with iLO5 and Silicon Root of Trust
- Maintaining system uptime and lowering exposure to security risks with fully integrated cluster software

HPE SOLUTION ENVIRONMENT AND CAE APPLICATIONS

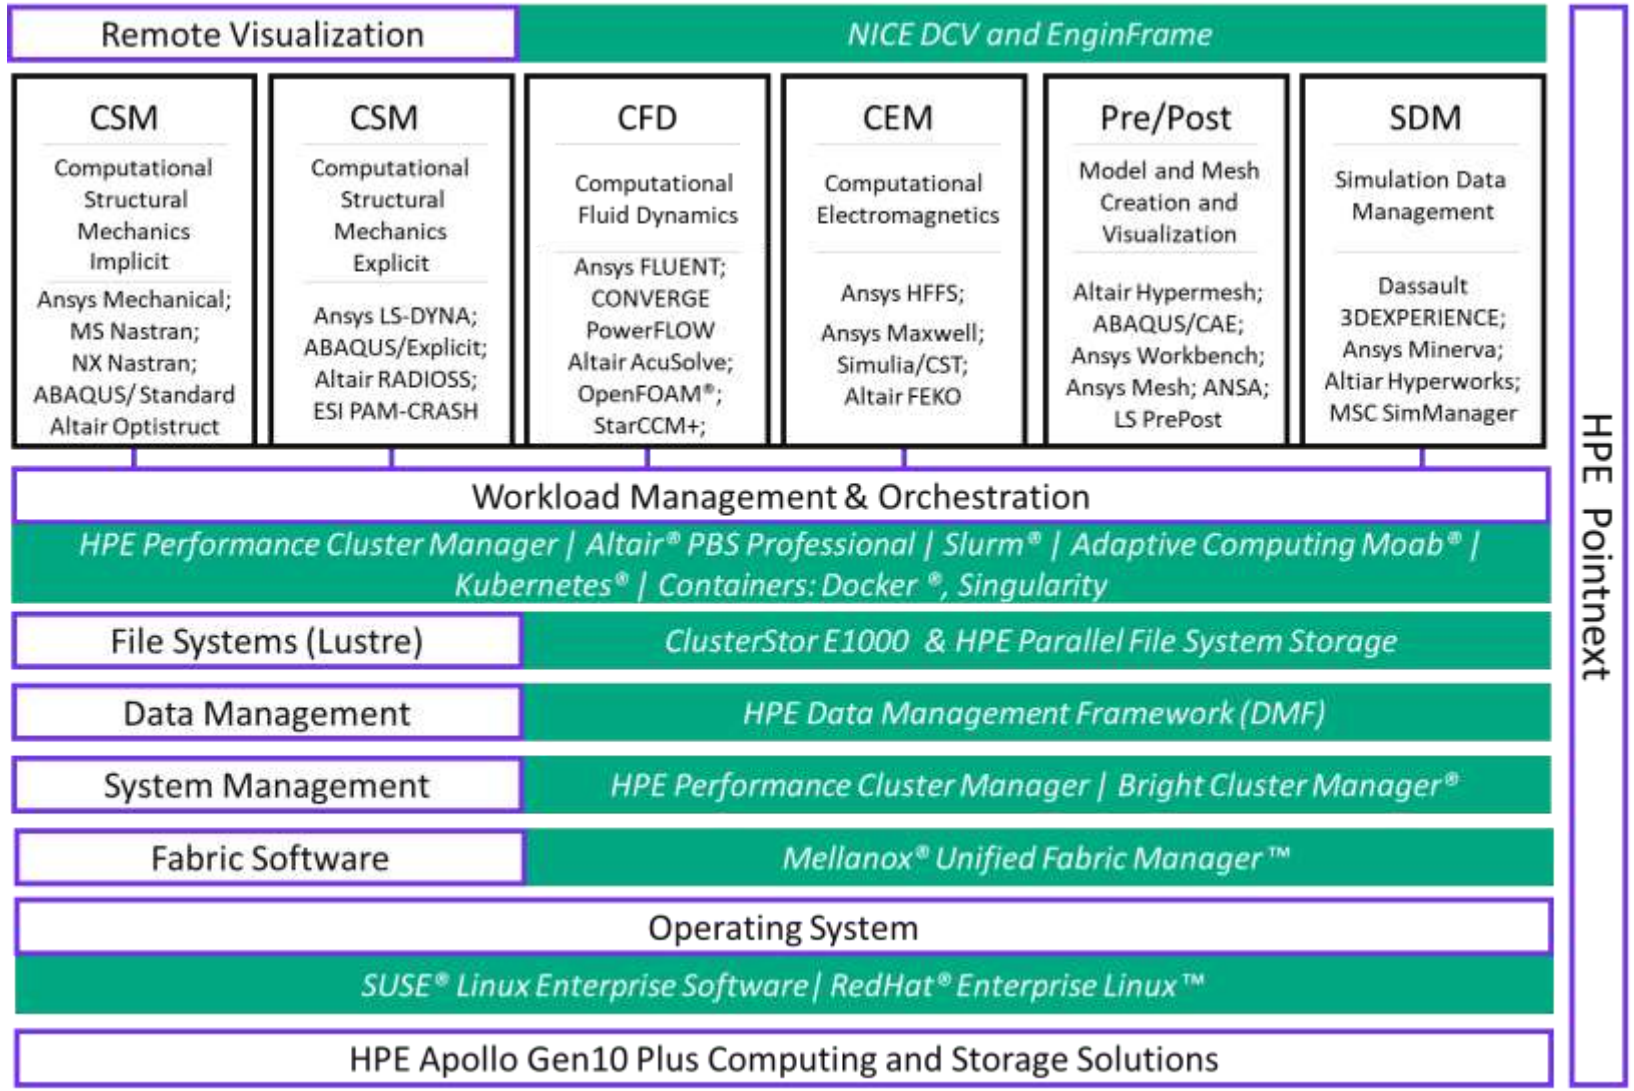

AMD EPYC Rome vs AMD Milan for CAE Workloads

AMD EPYC™ SOC ARCHITECTURE

2ND GEN EPYC VS. 3RD GEN EPYC COMPARISON – 9 DIE MCM (8 CCD + 1 I/O)

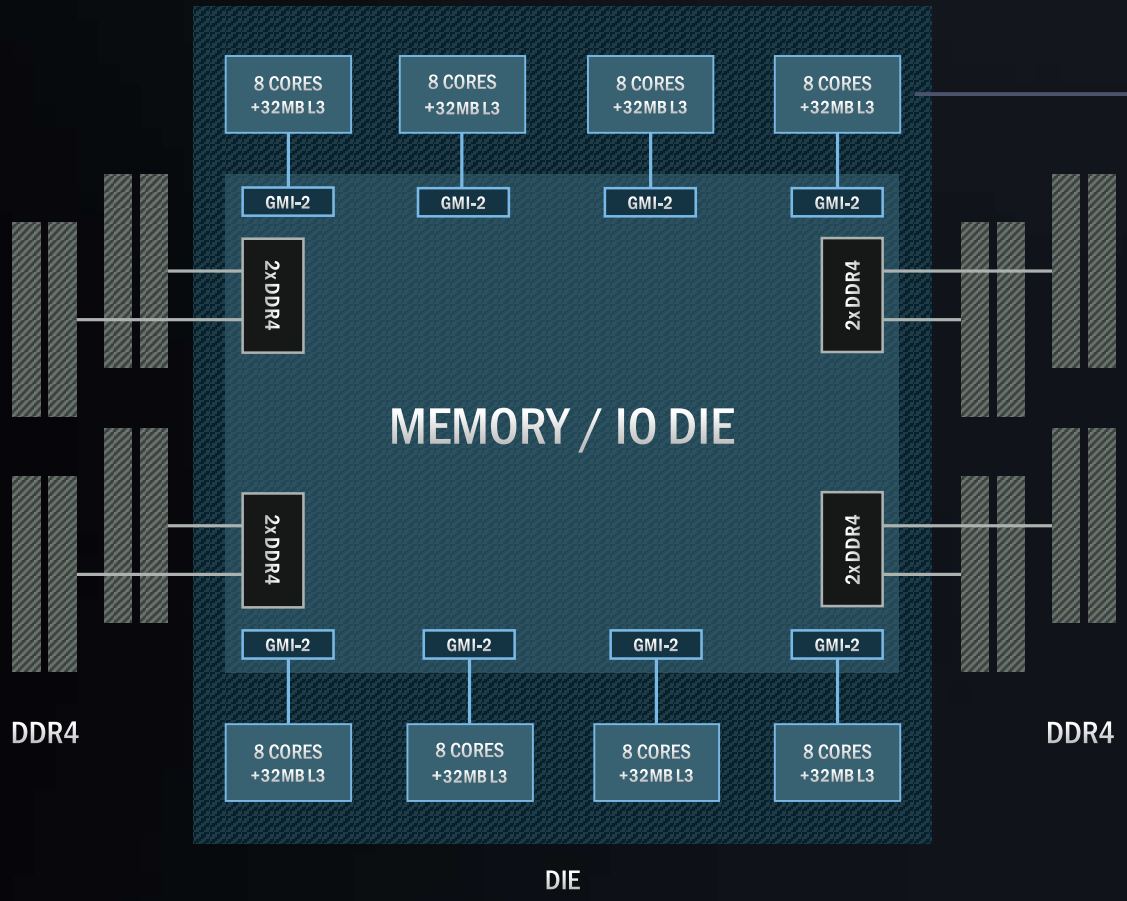

"ZEN 2"
AMD EPYC 7002

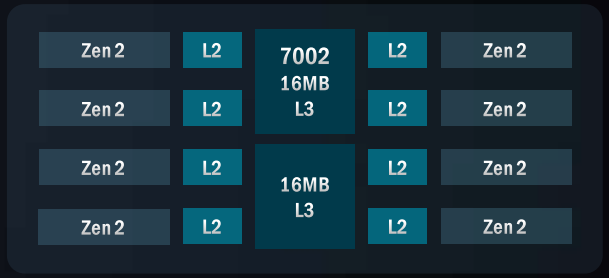

"ZEN 3"
AMD EPYC 7003

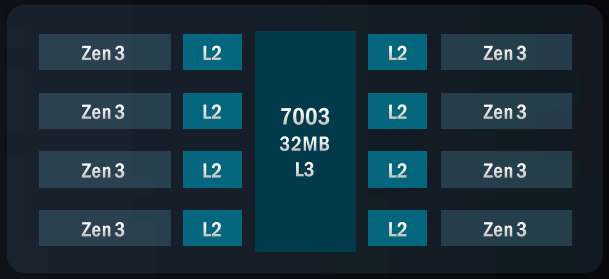

- Larger radix for 32MB of L3, enables better performance of large VMs/Cache sharing applications
- All 32MB can be allocated to single active core

EPYC™ 7543 SOC ARCHITECTURE (MILAN)

4 cores sharing 32 MB L3 Cache (8MB per core) and the 8 chiplets **provides efficient transfer of data** from cores to main memory for **higher memory bandwidth** to the 8 channels of DDR4 memory

EPYC™ 7532 SOC ARCHITECTURE (ROME)

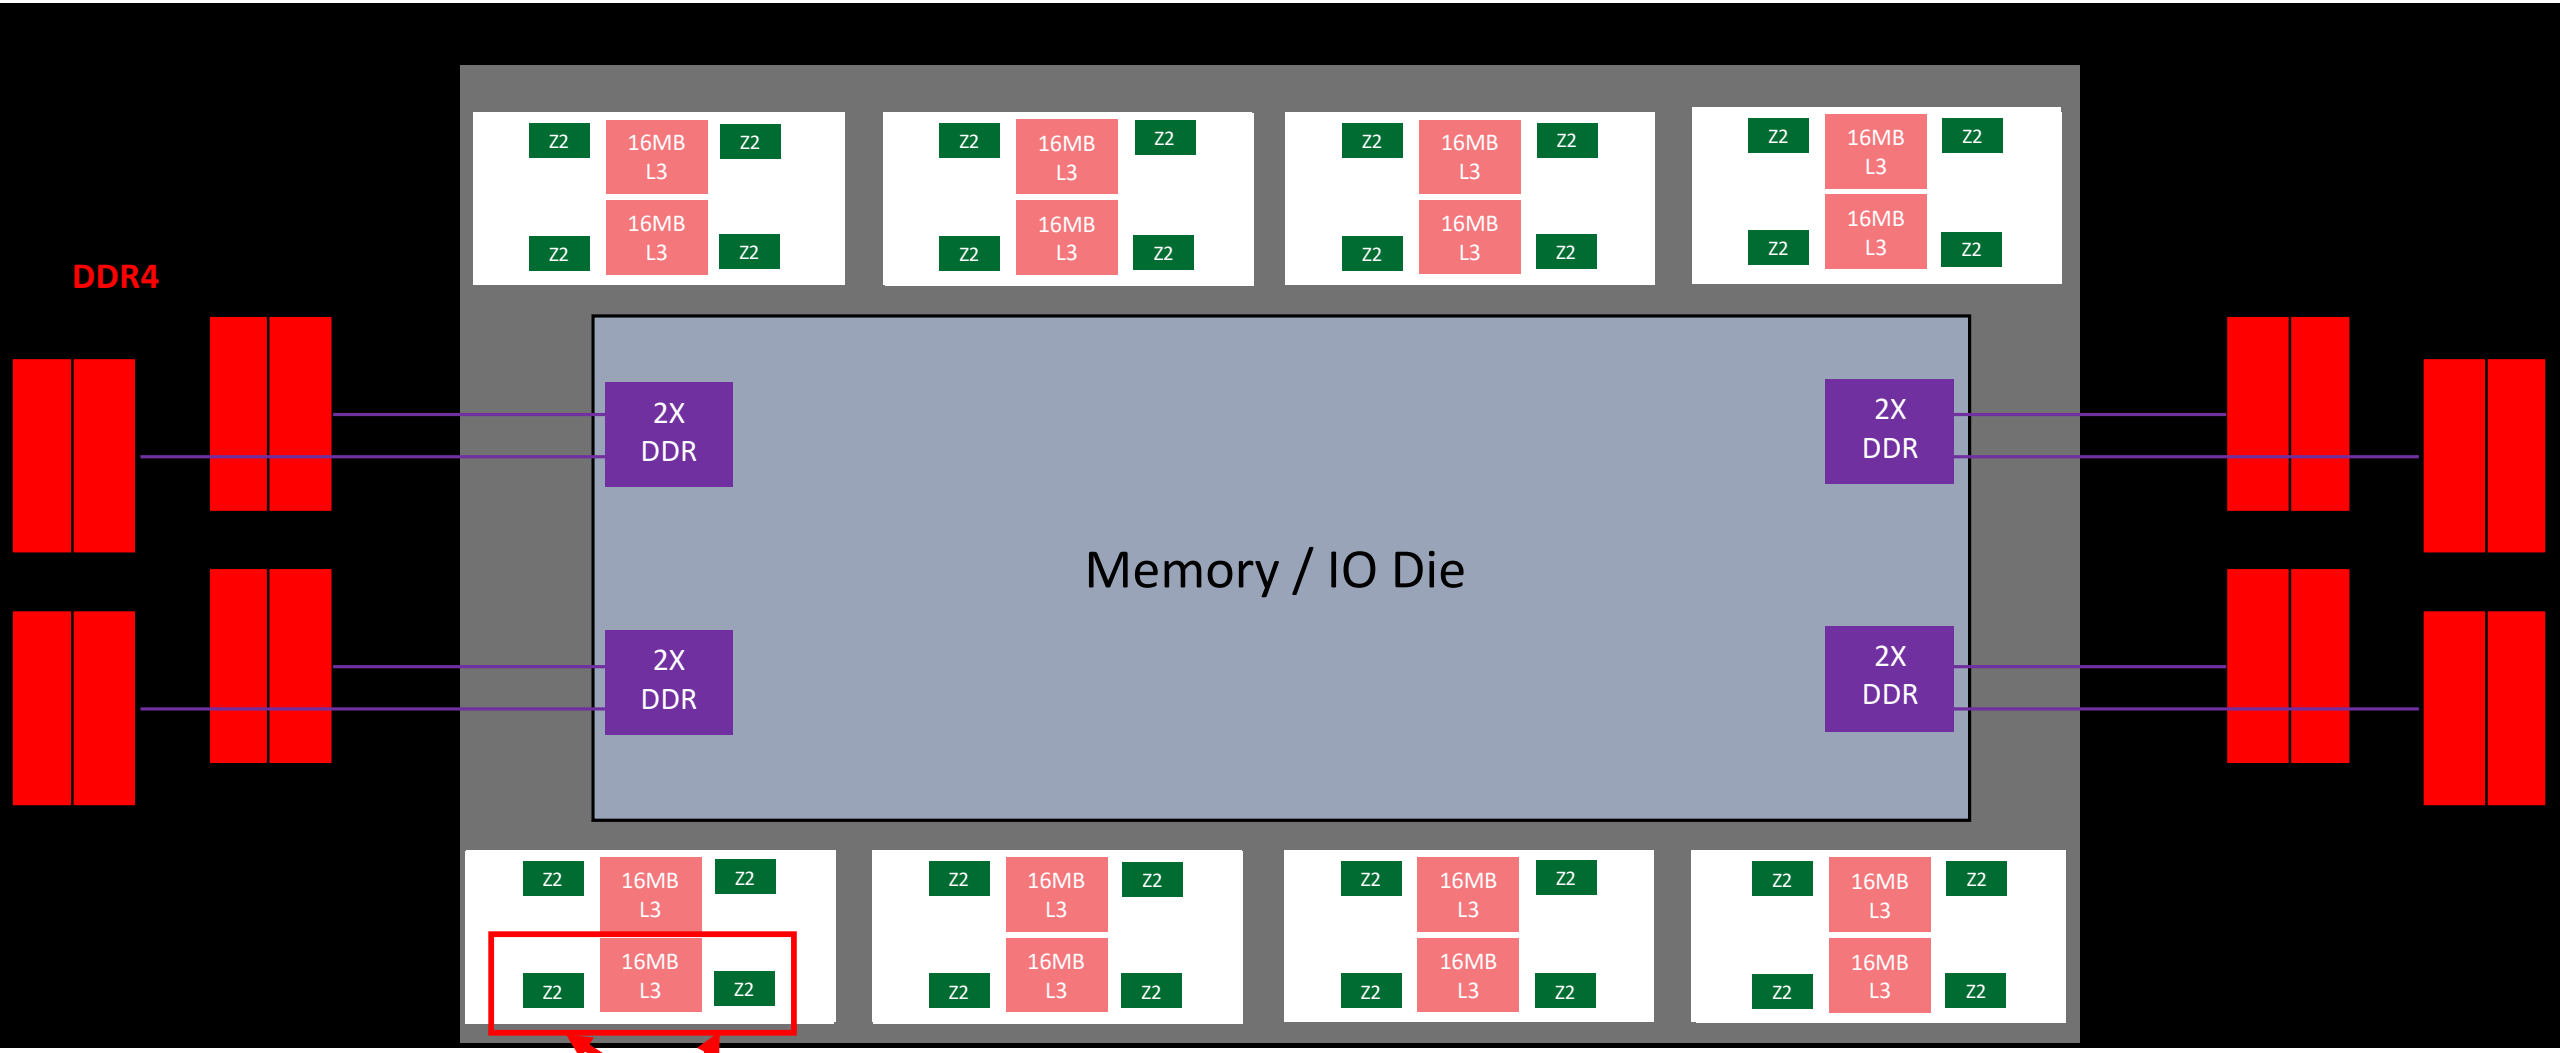

2 cores sharing 16MB L3 Cache (8MB per core) and the 8 chiplets providing efficient transfer of data from cores to main memory **for higher memory bandwidth** to the 8 channels of DDR4 memory

Ansys Standard benchmarks
AMD Rome 7532 32c vs AMD Milan 7543 32c

HPE BENCHMARK SYSTEM INFORMATION (AMD ROME VS AMD MILAN)

CPU	AMD® EPYC® ROME 7532	AMD® EPYC® Milan 7543
Base Frequency	2.4 GHz	2.8 GHz
Boost Frequency	3.3 GHz	3.7 GHz
Cores per node	32x 2 = 64	32x 2 = 64
L3/Last Level Cache	256 MB	256 MB
L3 Cache per core	8	8
Chiplets per CPU	8	8
Memory per node	256 GB	512 GB
TDP	200W	cTDP: 225-240W
Interconnect	InfiniBand HDR	InfiniBand HDR

ANSYS FLUENT STANDARD BENCHMARKS (JOBS PER DAY COMPARISON)

AMD EPYC 7532 2.4 GHz 32c vs AMD EPYC 7543 2.8 GHz 32c

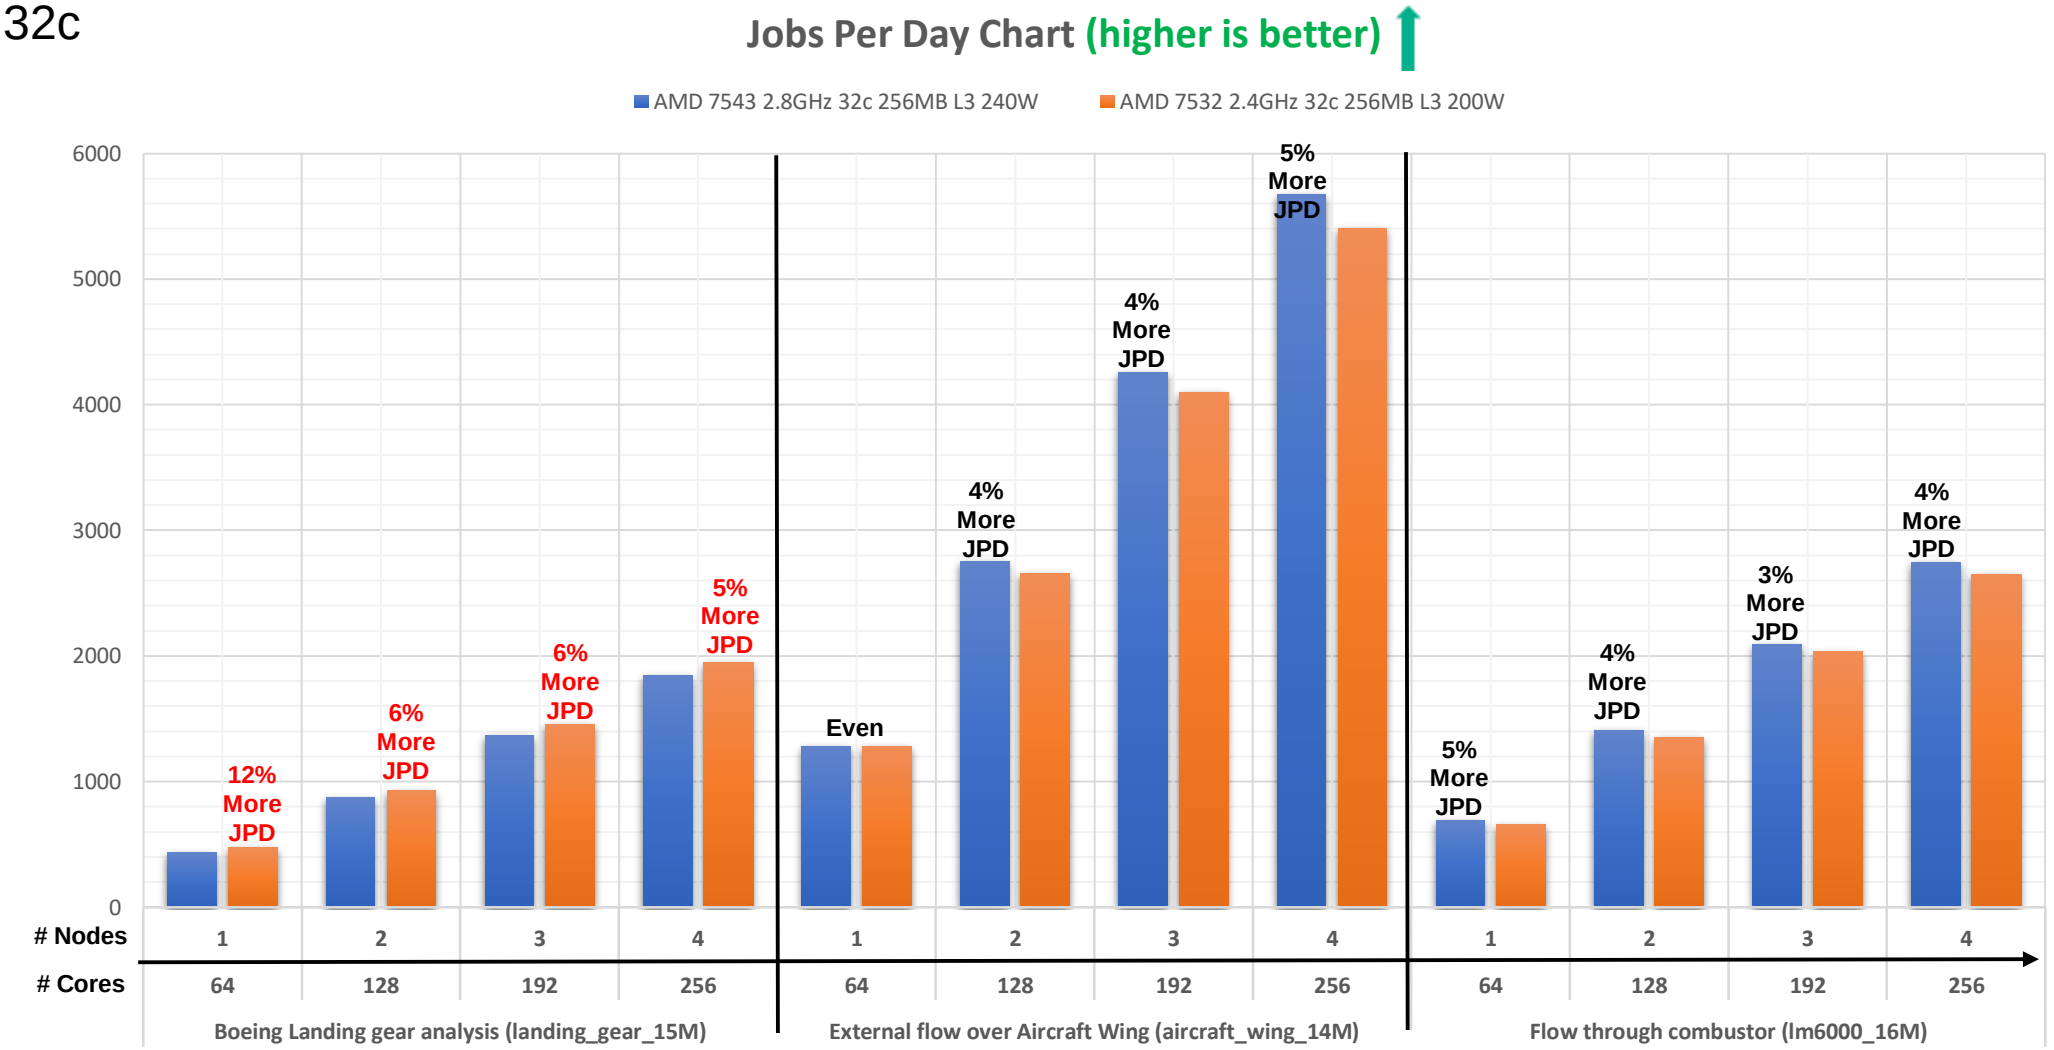

JPD= Jobs per day

ANSYS FLUENT STANDARD BENCHMARKS (WALLTIME)

AMD EPYC 7532 2.4 GHz 32c vs AMD EPYC 7543 2.8 GHz 32c

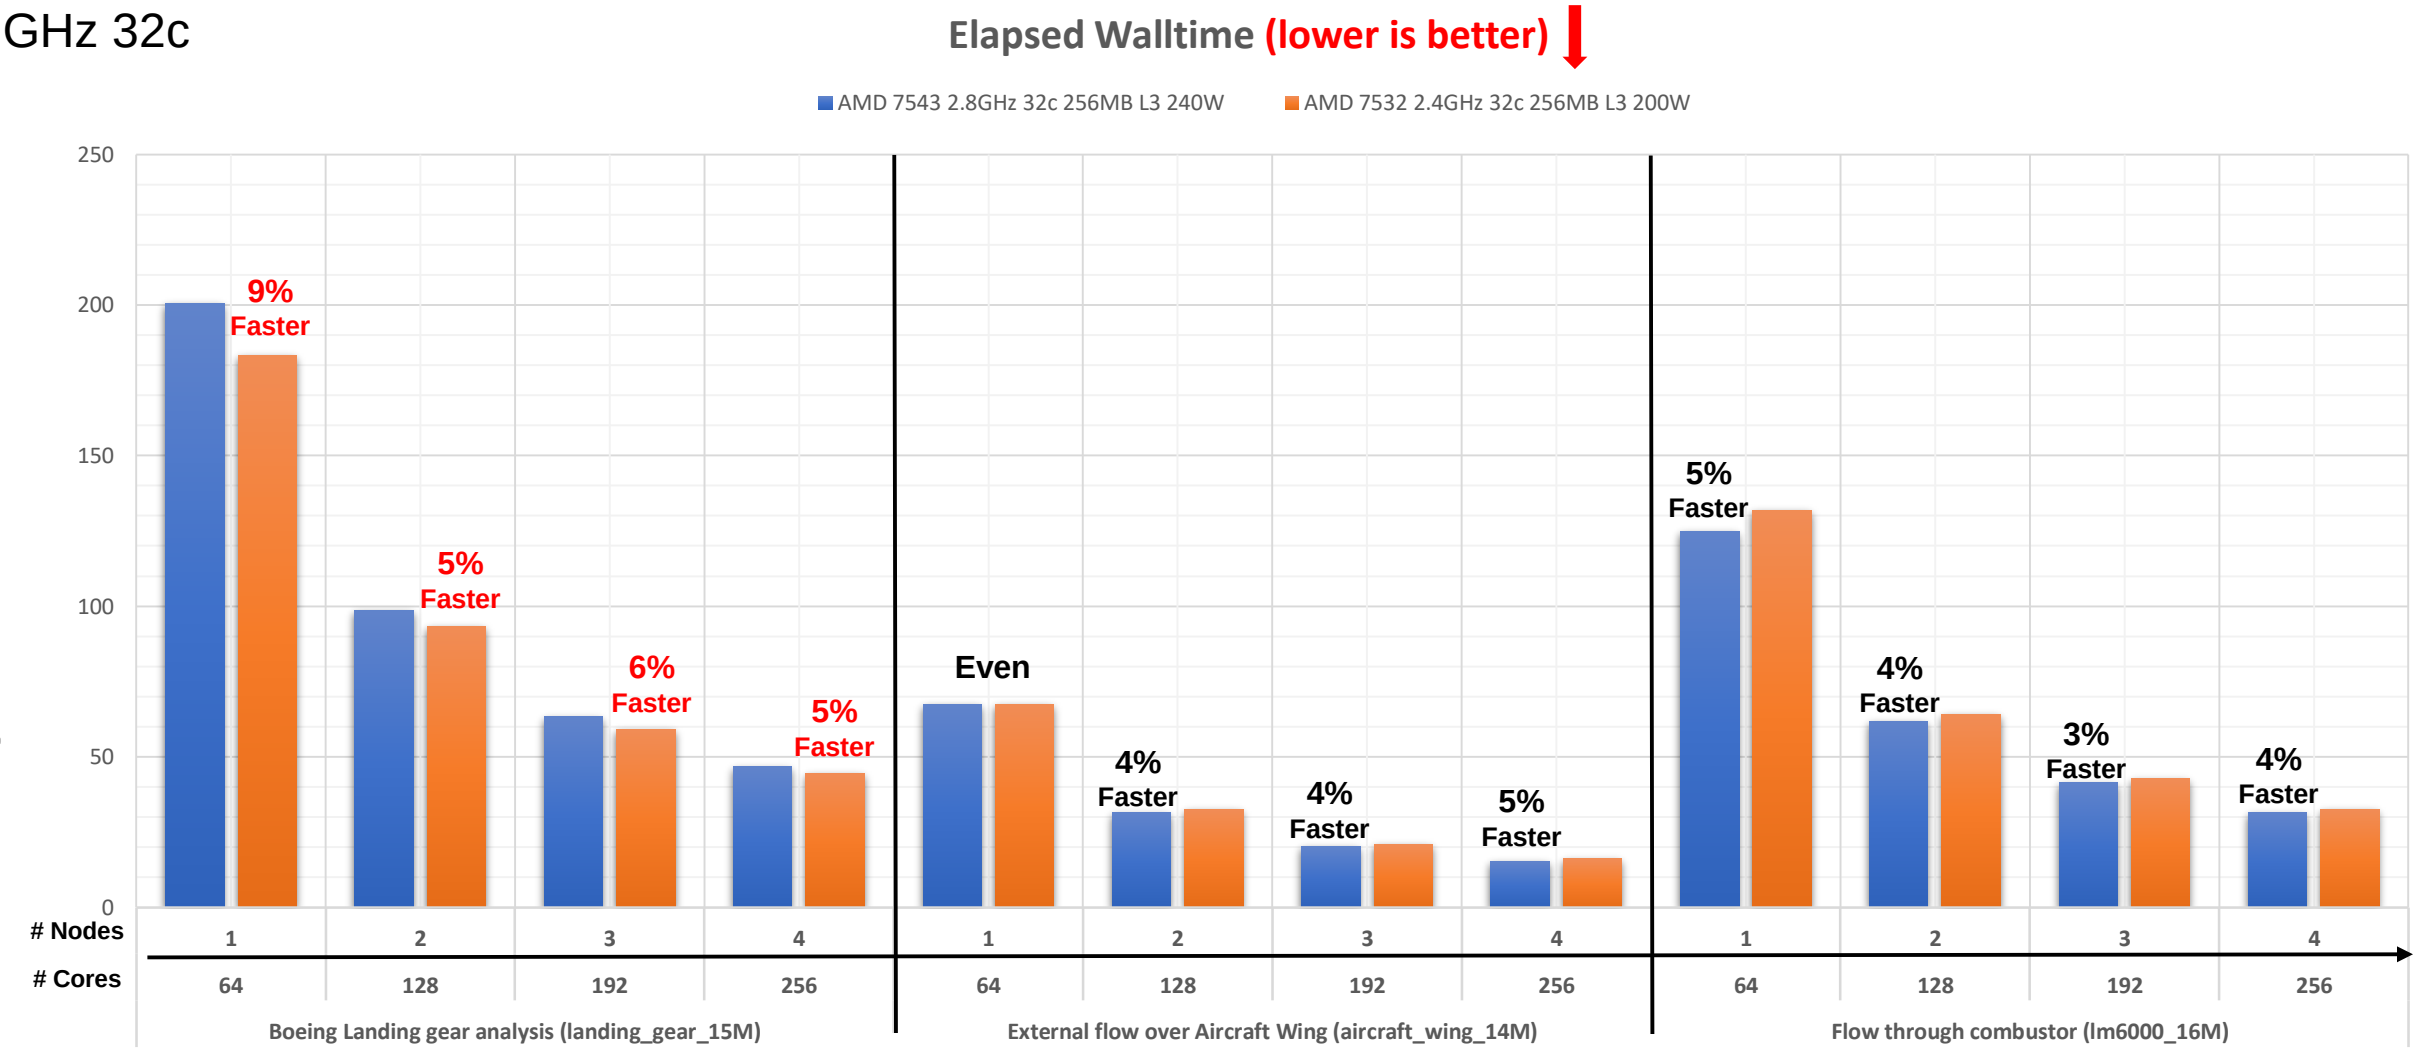

ANSYS FLUENT STANDARD BENCHMARKS (JOBS PER DAY COMPARISON)

AMD EPYC 7532 2.4 GHz 32c vs AMD EPYC 7543 2.8 GHz 32c

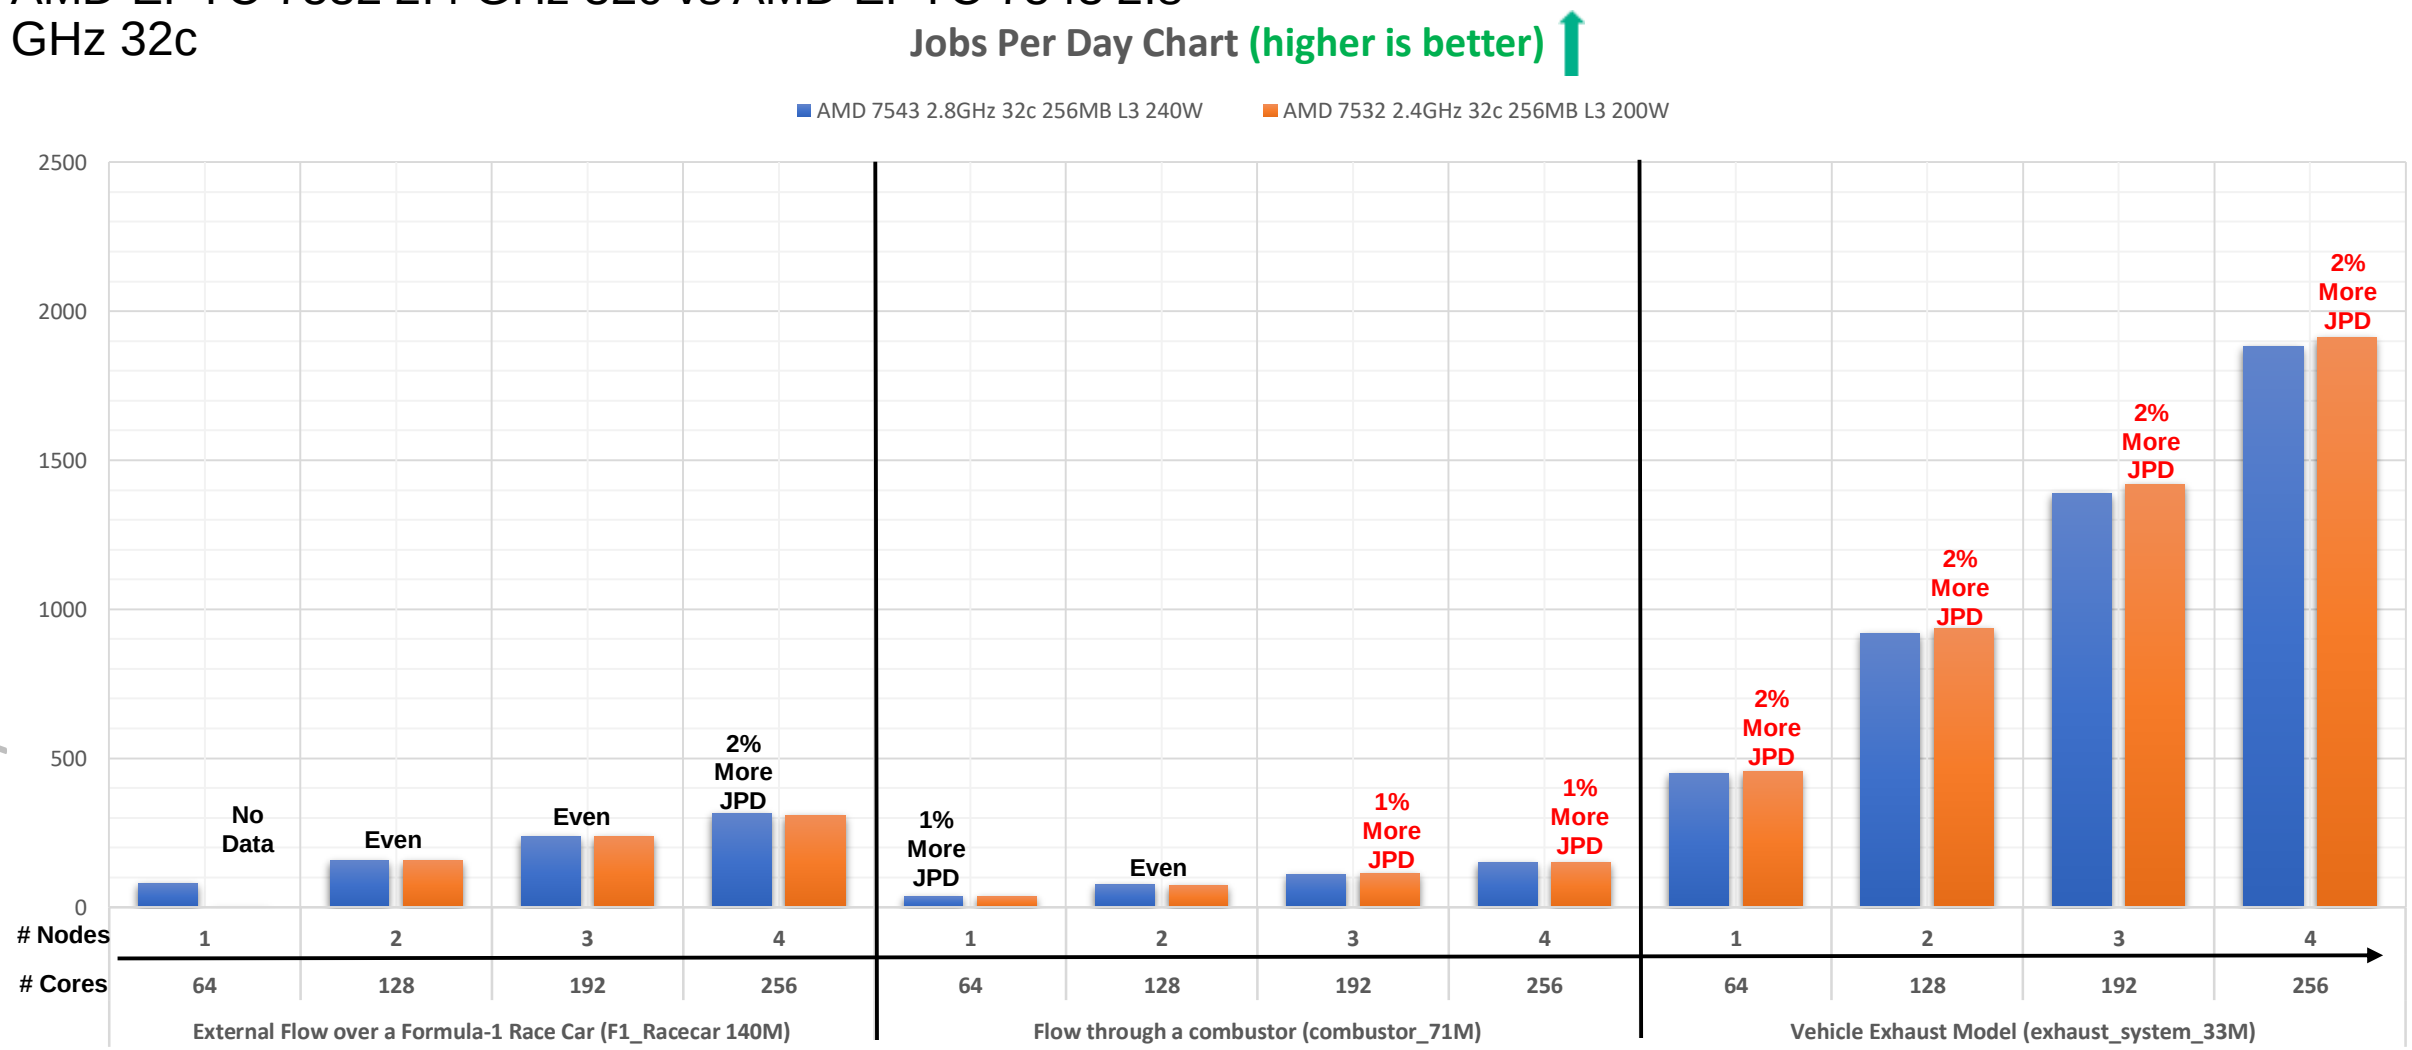

JPD= Jobs per day

ANSYS FLUENT STANDARD BENCHMARKS (WALLTIME)

AMD EPYC 7532 2.4 GHz 32c vs AMD EPYC 7543 2.8 GHz 32c

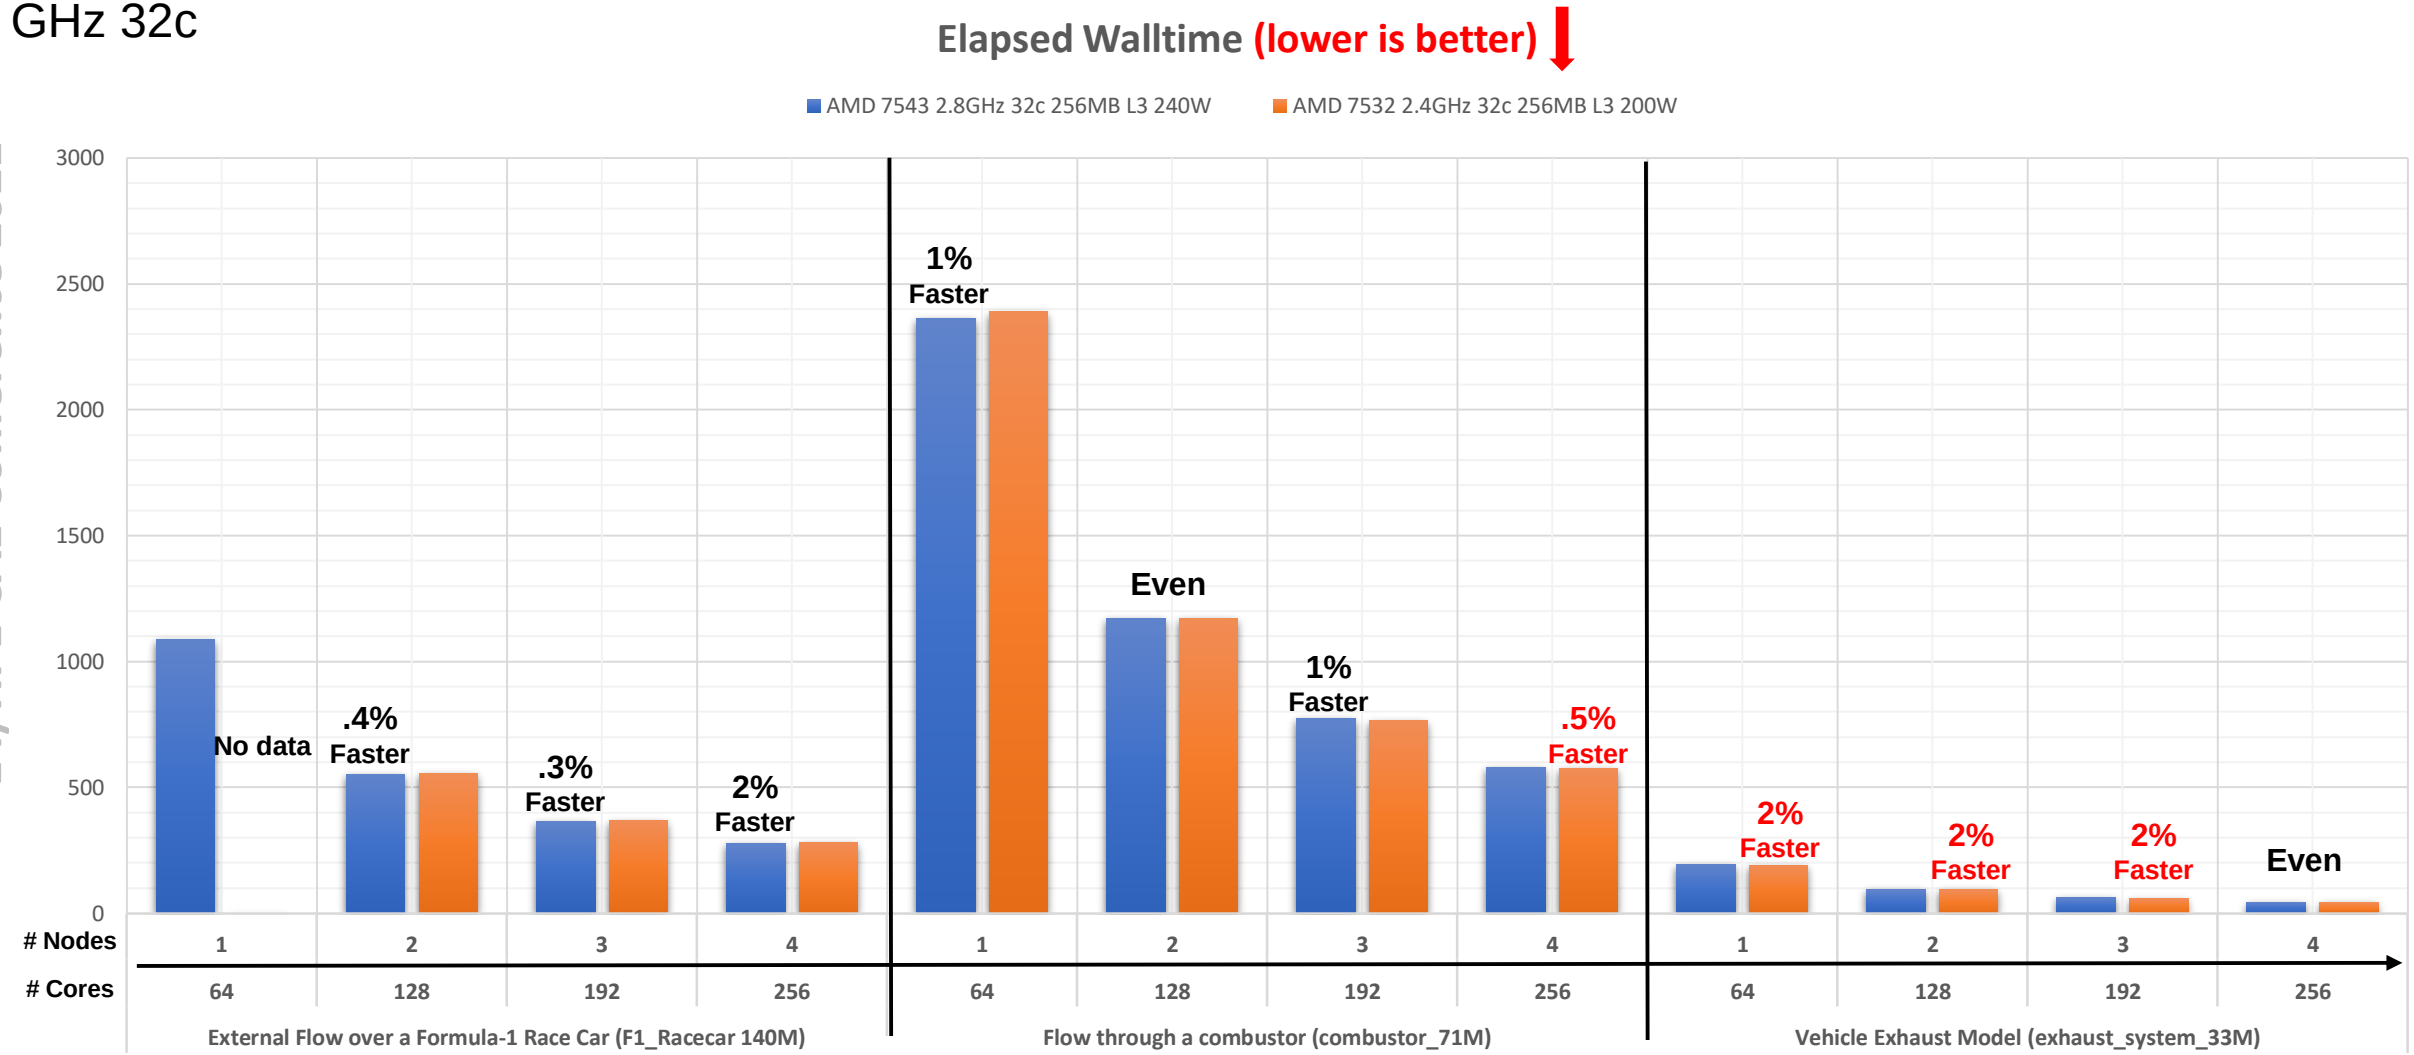

ANSYS CFX STANDARD BENCHMARKS

AMD EPYC 7532 2.4 GHz 32c vs AMD EPYC 7543 2.8 GHz 32c

Jobs Per Day Chart (higher is better)

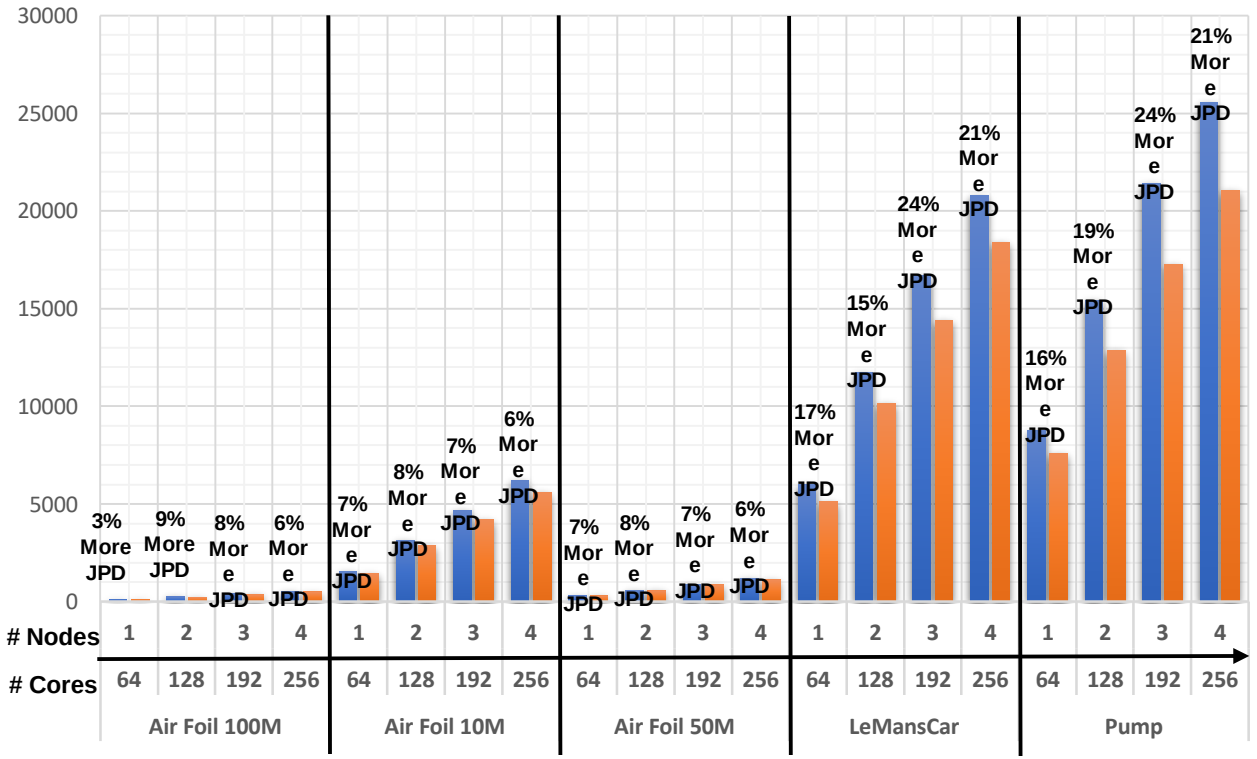

Elapsed Walltime (lower is better)

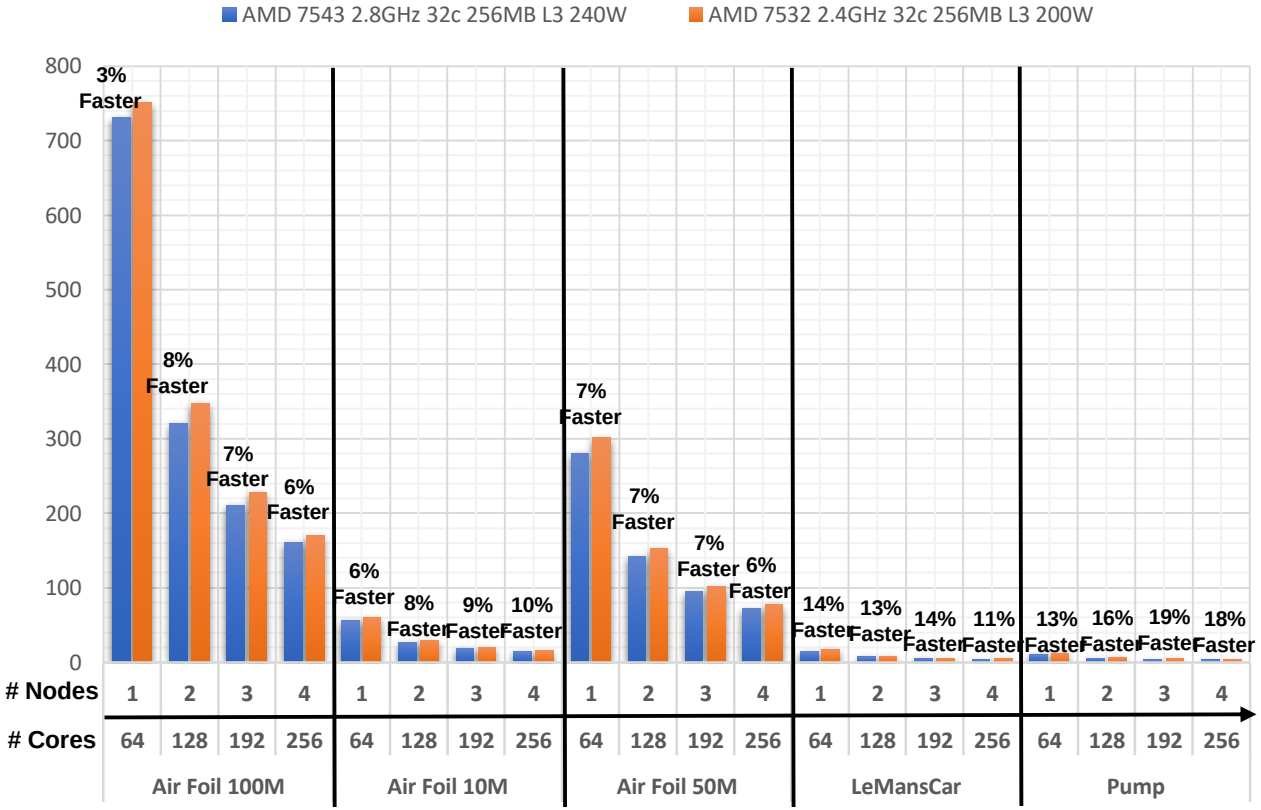

JPD= Jobs per day

ANSYS MECHANICAL STANDARD BENCHMARKS

AMD EPYC 7532 2.4 GHz 32c vs AMD EPYC 7543 2.8 GHz 32c

Jobs Per Day Chart (higher is better)

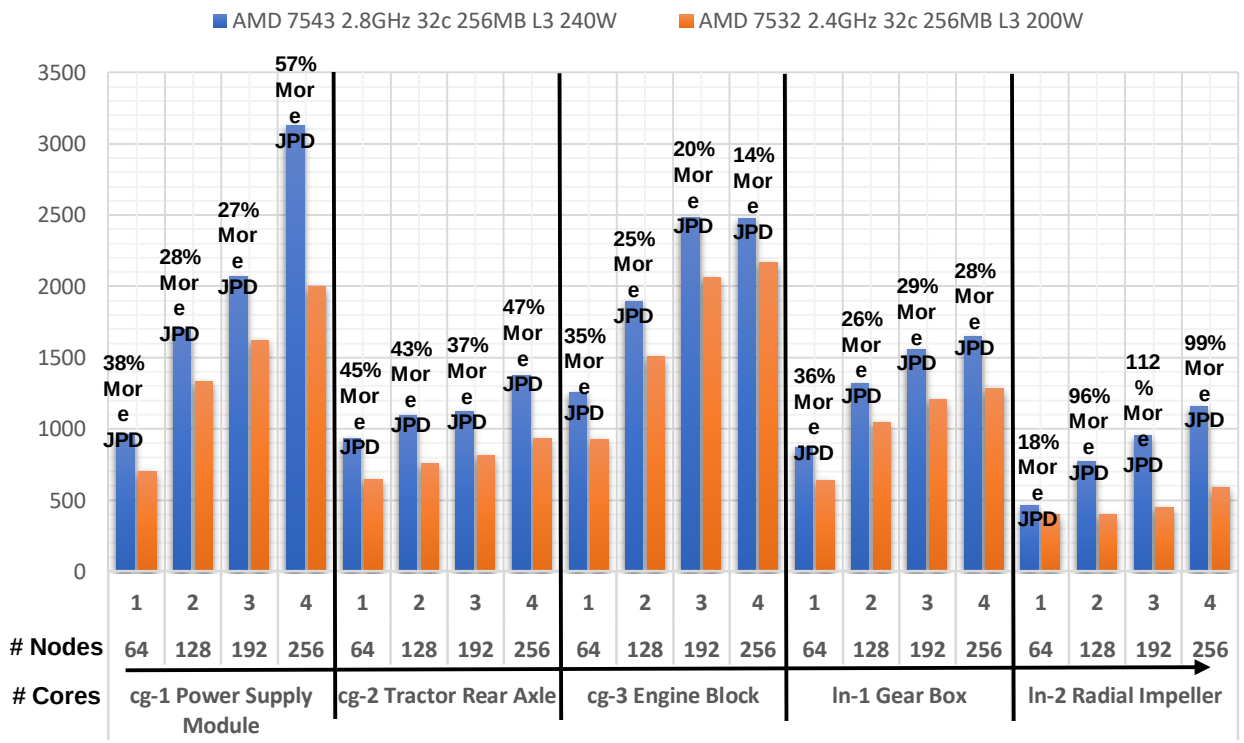

Elapsed Walltime (lower is better)

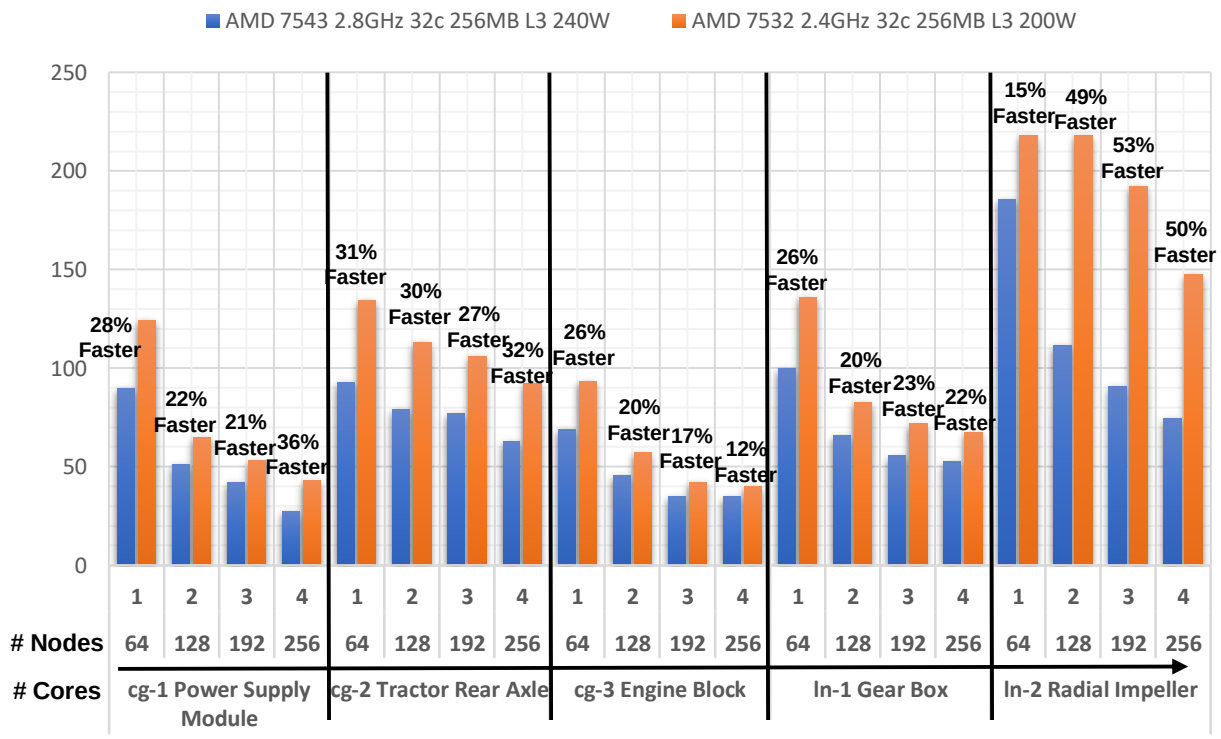

JPD= Jobs per day

The AMD 7543 (Milan) processor was run using Ansys Mechanical 2021 R2 that incorporates AMD math libraries which helped in the speedup over AMD 7432 (Rome) which was run with 2020 R2

ANSYS PLATFORM AND LINUX OS SUPPORT FOR MILAN VS ROME

AMD:

- To use AMD Milan SKU's the customer must use a cluster running either RHEL 8.3, CentOS 8.3 or SUSE 15 SP2 as their OS.
- If your customer needs to continue to run Ansys versions before 2021 R2 and they don't have an older cluster to run these simulations on then we recommend that the customer look at AMD Rome since AMD Rome is supported on RHEL 7.6 back to Ansys 2020 R1*
- The same goes with SUSE 15 SP2. If your customer wants to run older versions of Ansys than 2021 R2 and they don't have an older cluster to run these simulations on then we recommend that the customer look at AMD Rome since AMD Rome is supported on SUSE 12 SP4 back to Ansys 2020 R1*

*https://www.ansys.com/content/dam/it-solutions/platform-support/ansys-platform-support-strategy-plans-june-2021.pdf

HPE'S MATERIALS FOR CAE SOLUTIONS

Ansys HPE Solution Brochure with AMD

https://www.ansys.com/content/dam/company/technology-and-solution-partners/driving-innovation-ansys-hpe.pdf

Ansys HPE Solution Brief with AMD

https://www.ansys.com/content/dam/company/technology-and-solution-partners/enhancing-ansys-productivity.pdf

Business White Paper

https://psnow.ext.hpe.com/doc/a00021803ENW

Business White Paper w/AMD

https://psnow.ext.hpe.com/doc/a50002334ENW

Business White Paper Storage

https://psnow.ext.hpe.com/doc/a50003566enw

Technical White Paper

https://psnow.ext.hpe.com/doc/a00005355ENW

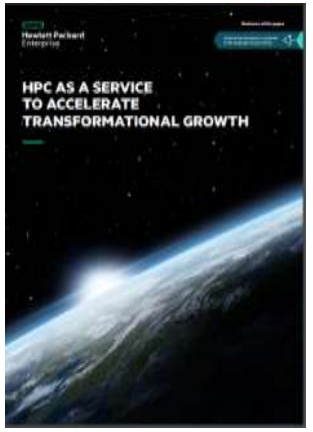

HPE HPCaaS White Paper

https://www.hpe.com/psnow/doc/a50004225enw

HPE GreenLake for HPC Brochu

https://www.hpe.com/psnow/doc/a50001688enw

HPE & E4: high tech and customer care

WHEN PERFORMANCE MATTERS

HPE and E4 Computer Engineering are long-time partners in the CAE market

HPE and E4 Computer Engineering jointly support users and enterprises in leveraging the benefits of Industry 4.0 via:

- Integration of AI/ML in the user's design workflows
- Integrating data (IoT) in the product 's development cycle
- Storing, managing and organizing large amount of data (IoT, legacy data repository)
- Co-designing highly scalable infrastructures, on-prem and cloud
- Benchmarking, performance evaluation, what-if analysis
- Infrastructure development plans

HPE & E4: Computational Intelligence and Deep Learning for Next-Generation Edge-Enabled Industrial Workflows

Tony De Varco (HPE)
anthony.devarco@hpe.com

Fabrizio Magugliani (E4 Computer Engineering)
fabrizio.magugliani@e4company.com
info@e4company.com

E4 COMPUTER ENGINEERING

Hewlett Packard Enterprise

Thanks!