

SLING



Dnevi SLING

Ansys HPC Solution

Marko Kolenc, Head of Office | Senior Application Engineer

Ansys delivers answers to today's toughest challenges

simtec

**ENGINEERING
WHAT'S
AHEAD**

ACCELERATE

*Speed Your Digital
Transformation*

INNOVATE


*Simulation Insights
Empower Optimization*

UNLOCK

*Realize the Potential in
Every Engineer*

Simulation is Essential for Digital Transformation

Simulation tools historically limited to design optimization used by specialists.



Ansys' capabilities extend simulation value across the entire product life cycle

SHIFT LEFT MORE INNOVATION FASTER CYCLE TIMES REDUCED COSTS AND RISKS **80%** OF DEVELOPMENT COSTS ARE LOCKED IN EARLY IN THE CONCEPT AND DESIGN PHASES

9X ↓
REDUCTION IN DEVELOPMENT TIME

1,000X ↓
REDUCTION IN VALIDATION TIME FOR COMPLEX SYSTEMS

25% ↓
WEIGHT REDUCTION WITH TOPOLOGY OPTIMIZATION ADDITIVE MANUFACTURING

10-20% ↓
REDUCTION IN O&M COSTS

Reduction numbers based on industry-specific customer testimonials

50 Years of Simulation Innovation and Leadership

simtec

50
YEARS OF
INNOVATION



Shatter Records



Unlock Possibilities



Make the Unmakeable



Save Lives

#1
IN
SIMULATION

GLOBAL

PROVEN

FOCUSED

COMMITTED

6,000+ Full-Time Employees
90+ Offices Worldwide

\$2.07 Billion USD
2022 Revenue

Simulation is
All We Do

93% 2022 Customer
Satisfaction Score



ANSYS

ANSYS ongoing investment in critical simulation capabilities **simtec**

29

STRATEGIC ACQUISITIONS

\$5B+

INVESTED IN ACQUISITIONS

350+

TECHNOLOGY PARTNERS

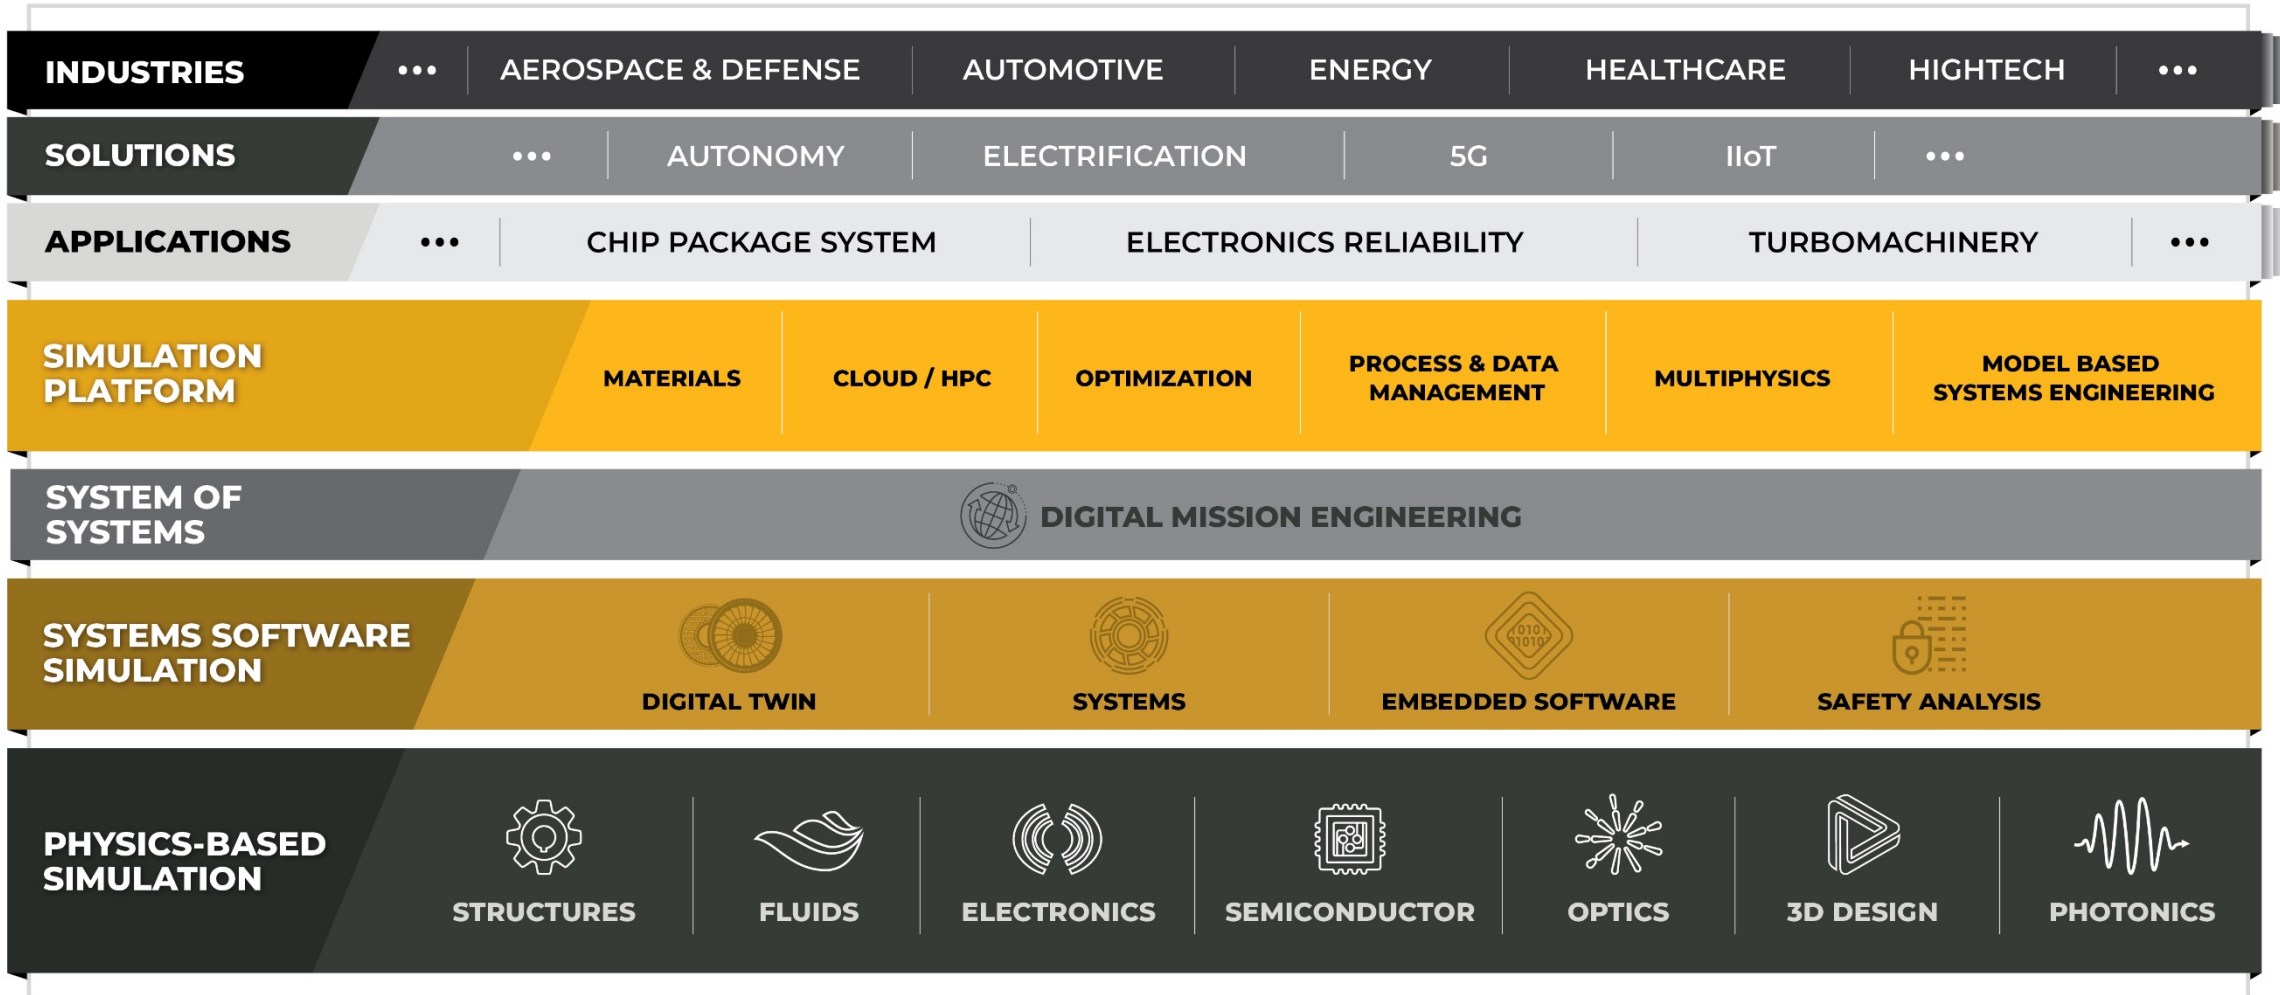
Organic software development and improvement with revolutionary advancements in physics and algorithms.

Strategic acquisitions to continuously build the platform and enhance our staff.

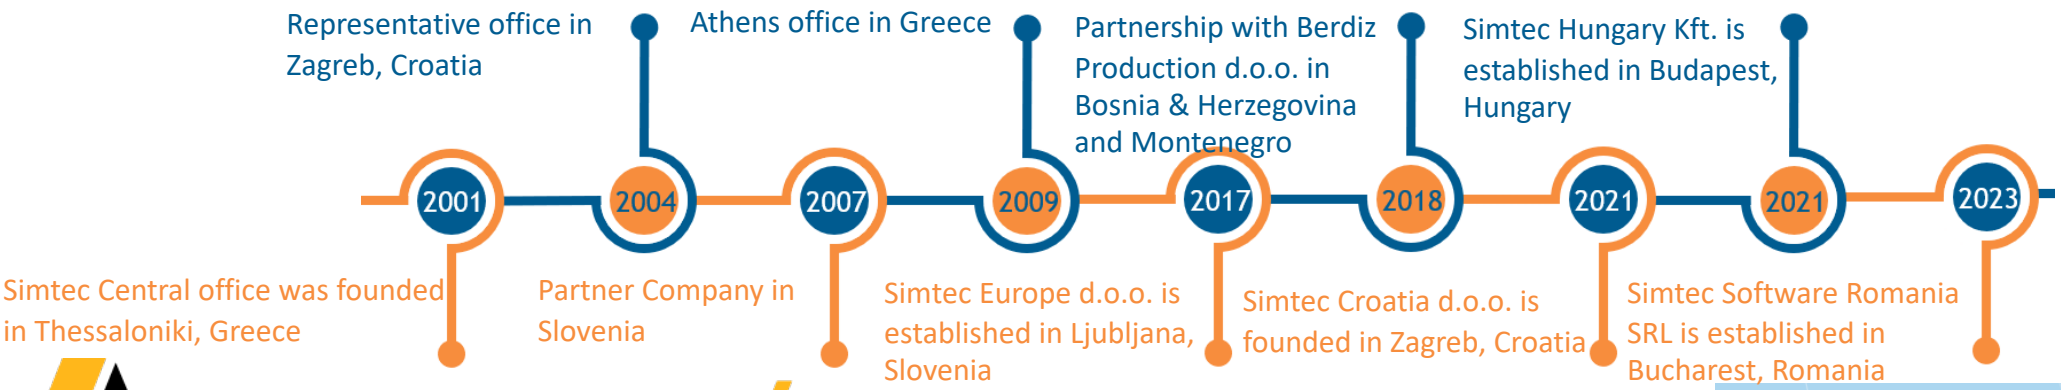
Strategic partnerships to expand our reach and multiply our impact.



ANSYS Comprehensive Capabilities Drive Value Creation



Simtec – Ansys Elite Channel Partner



ELITE CHANNEL PARTNER



Distribution

- Sales
- Support

Training

- Beginner
- Advanced

Consultation

- Project
- Workflow



Our customers




HPC Strategy

What does HPC give the end user?




Ansys Fluent & HPC

- Higher mesh counts and greater mesh density
- Longer Transients
- More complex physics



Ansys Mechanical & HPC

- Run simultaneous **design points**
- Run **larger models** that require **more memory**



Ansys HFSS & HPC

- Run multiple **frequency sweeps in parallel**
- Get **enough memory** to run large models

*Run
multiple
designs in
parallel*

*Reduce
solve
time*

HPC development enables bigger, better and faster simulations

Maximizing innovation through 3 strategic HPC focus areas



HPC Development

Intense focus on delivering HPC performance and capability provides customers with higher levels of simulation fidelity, engineering insight and innovation.

Parametric parallel, domain decomposition in HFSS reduces time to solution

- HFSS Frequency sweep extraction leverages HPC
- Frequency points solved in parallel
- Leverages both shared and distributed (elastic) hardware
- Fully automated, no user intervention required

Parallel processing on GPUs in HFSS SBR+ enables significant performance increase

- HFSS Mesh Fusion
- Breaks problem
- Allows noncon
- Maintains HFSS


Real-Time Radar is a subset of the HFSS SBR+ solver adapted to run on GPUs

- Over 3000x perform
- CPU implementation
- Able to generate m
- imagery for smart i
- (milliseconds/frame
- New application: si
- gnal propagation, i
- nterference in cor

How hybrid parallel improves Mechanical scaling at high core counts



Enhanced domain decomposition scales LS-DYNA to very high core counts




HPC Partnerships

Collaborations with HPC leaders enable breakthrough productivity on current and emerging HPC infrastructure – on premises and in the cloud.


HPC partnership with Intel: oneMKL and Fluent collaboration results in performance improvement

- Sparse LDU smoother introduced follows Inspector-Executor approach and thus enables matrix preparation Intel Advanced Vect
- Ansys Fluent demo performance impro
- Ansys and Intel con
- fine-tune performa



HPC partnership with AMD: Integration of optimized libraries in Mechanical results in up to 2.1X speedup


Performance of Ansys Mechanical 2021 R1 vs. 2021 R2 (w/ optimized AMD AOCL BLUS) on AMD EPYC 73F3 (16 Cores)



HPC partnership with ARM/AWS: porting of Fluent on ARM yields compelling performance / cost ratio on Graviton3E



HPC partnership with NVIDIA: new fully native multi-GPU solver in Fluent shows astounding performance (per watt) gains




HPC Deployments

Offer cost-effective HPC on-premises and in the cloud, addressing customers' large variety of licensing and compute needs, from average to burst to fluctuating workloads.

HPC Deployments: Ansys supports the use of HPC resources anywhere

- For on-premises, Ansys tests, supports and certifies:
 - Leading remote display software solutions (VNC, DCV, Exceed onDemand, Exceed TurboX and Windows RDP)
 - 4 proven VDI solutions
 - Leading job sc


HPC Deployments: Ansys provides scalable, value-based HPC licensing



HPC Deployments: Ansys provides the best ROI of on-premises HPC



HPC Deployments: Ansys offers a choice to scale HPC to cloud



Ecosystem of Ansys partners providing HPC solutions & services



HPC partners

- ✓ Recommend and/or deliver high-performance, **Ansys-validated HPC systems** with wrap-around support, ensuring quick deployment, and performance benchmarking

Hardware partners

- ✓ Deliver **hardware optimized** for Ansys applications with fast processors, ample memory, and high-end graphics, along with Ansys performance data



Cloud partners

- ✓ Enable **rapid simulation access** with BYOC options for quick deployment of Ansys applications in users' clouds, using existing SLAs

System integrators

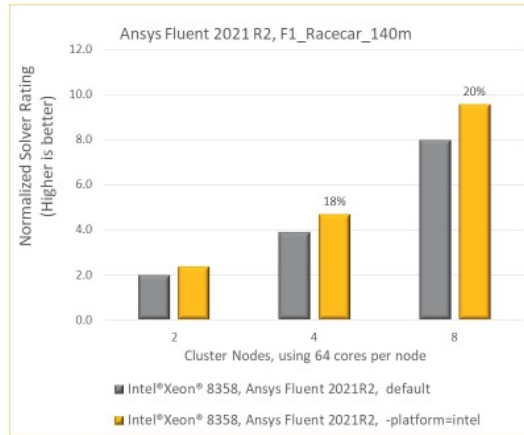
- ✓ Offer **full-service support**, incl. design, deployment, management, and HPC admin tasks, with integration of on-premises and cloud solutions



HPC partnerships are key to support our R&D


HPC partnership with Intel: oneMKL and Fluent collaboration results in performance improvement

- Sparse LDU smoother introduced follows Inspector-Executor approach and thus enables matrix preparation in a form more suitable for Intel Advanced Vector Extensions
- Ansys Fluent demonstrates up to 20% performance improvement using LDU smoother
- Ansys and Intel continue to work together to fine-tune performance of sparse smoothers





Performance varies by use, configuration and other factors. See endnotes for configuration details.


HPC partnership with AMD: Integration of optimized libraries in Mechanical results in up to 2.1X speedup




HPC partnership with NVIDIA: new fully native multi-GPU solver in Fluent shows astounding performance (per watt) gains



HPC partnership with ARM/AWS: porting of Fluent on ARM yields compelling performance / cost ratio on Graviton3E



performance for Graviton3E is ~63% better than for Graviton2



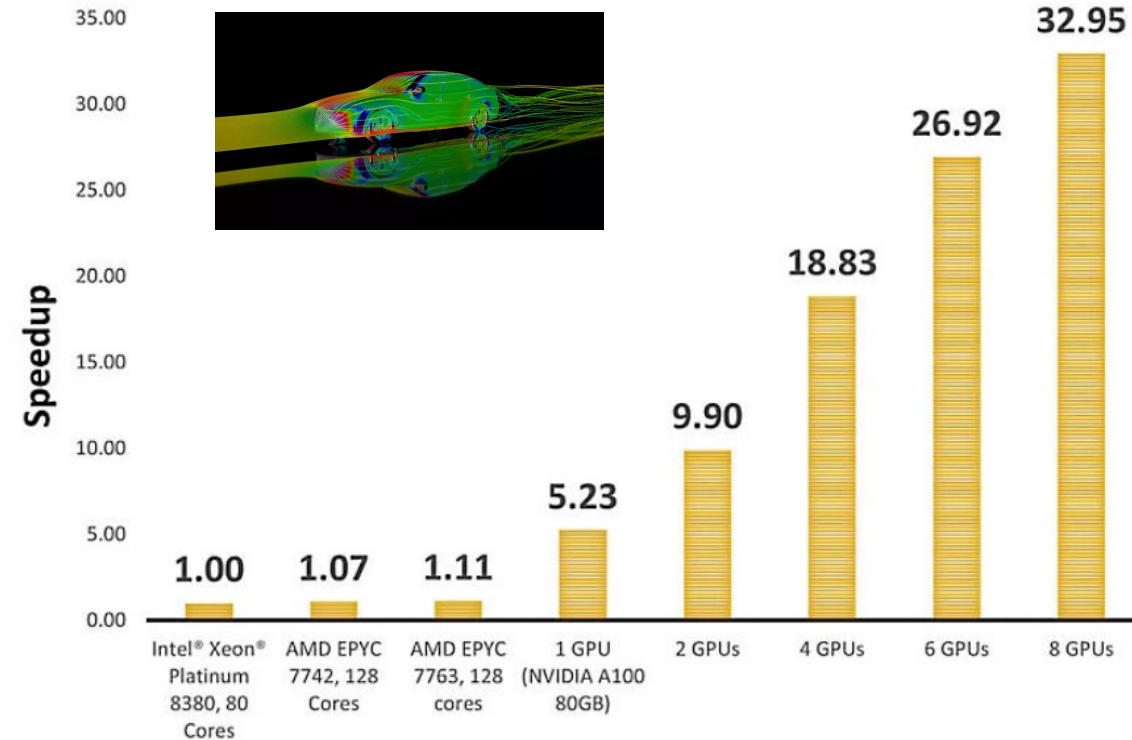
cost per simulation is ~2.7x better on Graviton3E compared to Graviton2

Ansys Commitments

- Fully native multi-GPU solver in **Fluent 2023** shows that **4 high-end GPUs provide the same performance as more than 1,000 CPUs**
- **Ansys Mechanical 2023** demonstrates up to **43% performance gains on 4th generation of Intel Xeon Scalable processors vs. previous generation (due to AVX-512 support)**.
- **Ansys LS-DYNA** is up to **32% faster on 4th Gen Intel Xeon Scalable processors from Intel AVX-512 relative to SSE2 technology on the same hardware**.

Customer Benefits

- Better utilization of hardware, software licenses, and people.
- More time for analyzing results, leading to improved product understanding.
- Accelerate simulation throughput.
- Get more insight into product performance.
- Evaluate more product design ideas

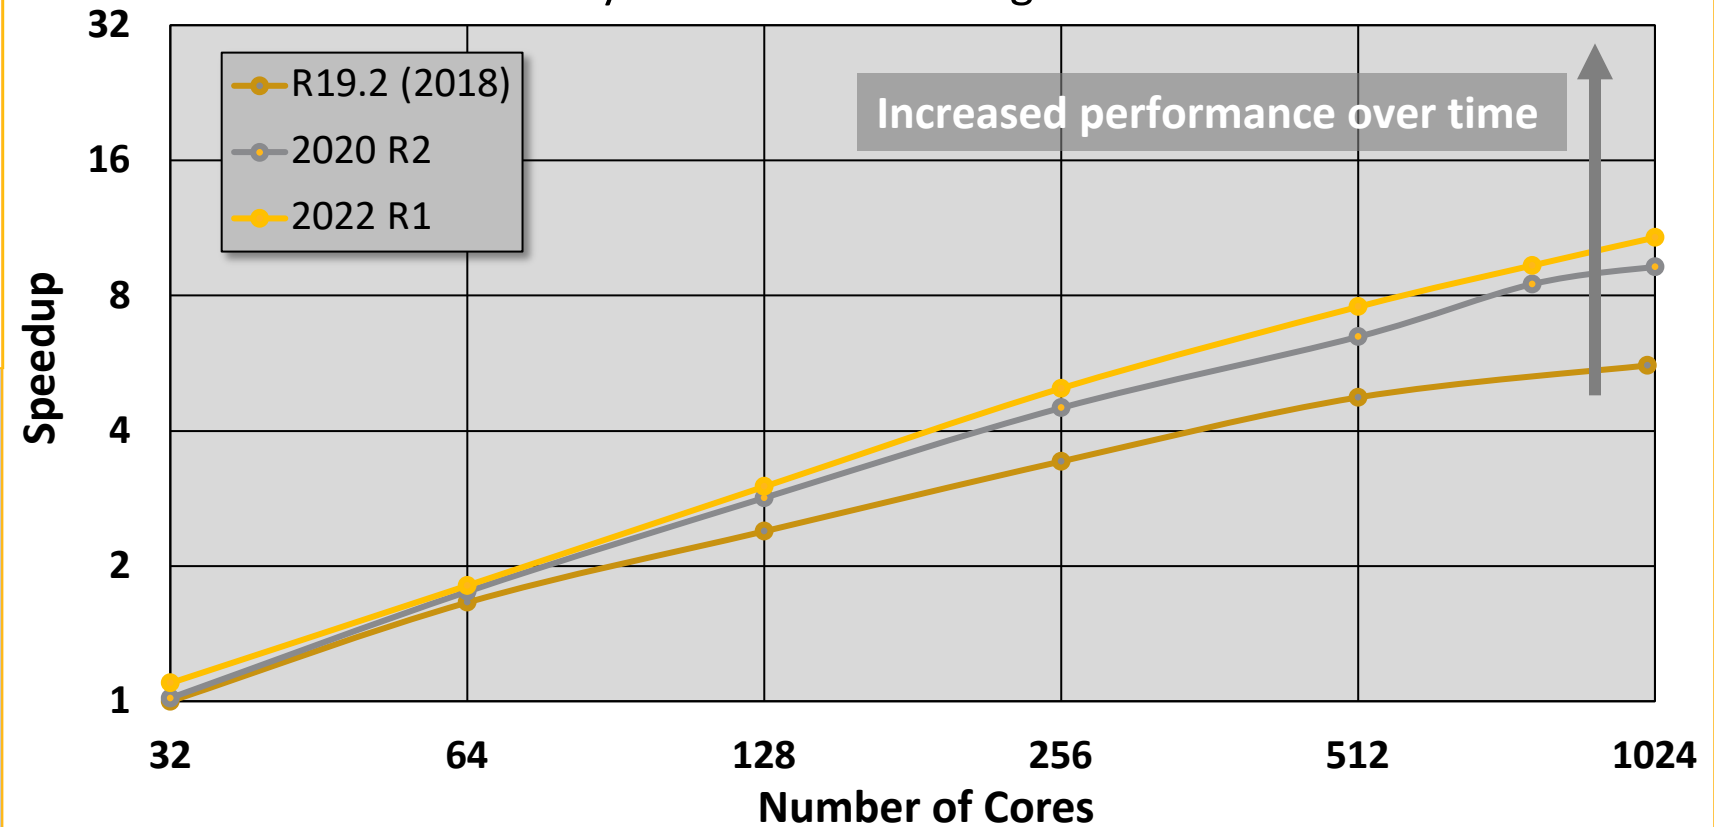


Example of Ansys Mechanical




- 5.6 million DOF; sparse solver
- Nonlinear static analysis involving contact, constraint equations, unsymmetric matrices
- Linux cluster; each compute node contains 2 Intel Xeon Gold 6148 processors (40 cores), 384GB RAM, SSD, Mellanox InfiniBand, CentOS 7.6

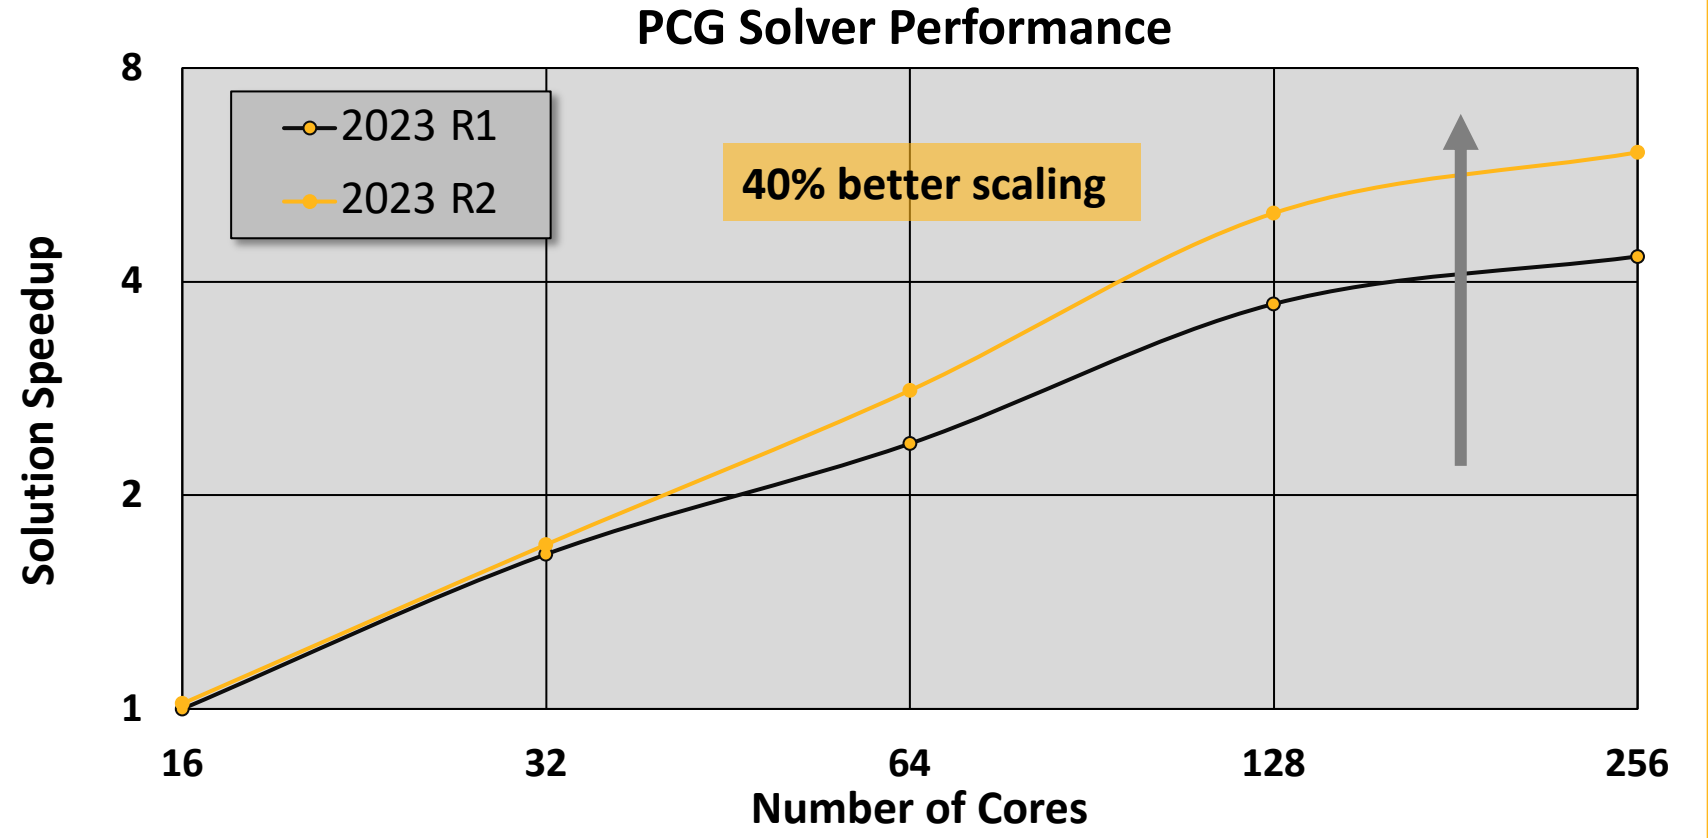
Ansys Mechanical Scaling Performance



Example of Ansys Mechanical



- 63 million DOF
- Static analysis with SOLID186/187 elements, PRET179 elements, penalty contact, 47000 constraint equations
- Linux cluster; each compute node contains 2 Intel Xeon Gold 6142 processors, 384GB RAM, SSD, CentOS 7.9, Mellanox InfiniBand




Example of Ansys Fluent

Sliding mesh capability

- Enhanced mesh intersection algorithms
- Incremental neighborhood to reduce mesh update overhead
- Enhanced wall distance calculation
- Enhanced data exchange strategies



F1_SM_140M with rotating wheels



NextGen Native GPU CFD Solver






GPU: What's the Excitement About?




Engineers want to use simulation to see how their ideas will perform against millions of variables

- Huge numbers of simulations
- Reduce time to market
- Need drastic increase simulation throughput


→ Without compromising accuracy!




The time is now for general purpose GPUs in scientific computing




- ANSYS Flagship Products such as Fluent, Mechanical, HFSS and Maxwell have used GPUs as accelerators in “offload” mode for years
 - GPUs have large bandwidth and computational throughput
 - Good for isolated, expensive computations
 - Solver modules like linear algebra, ray tracing, radiation models can run efficiently on GPUs



Fluent




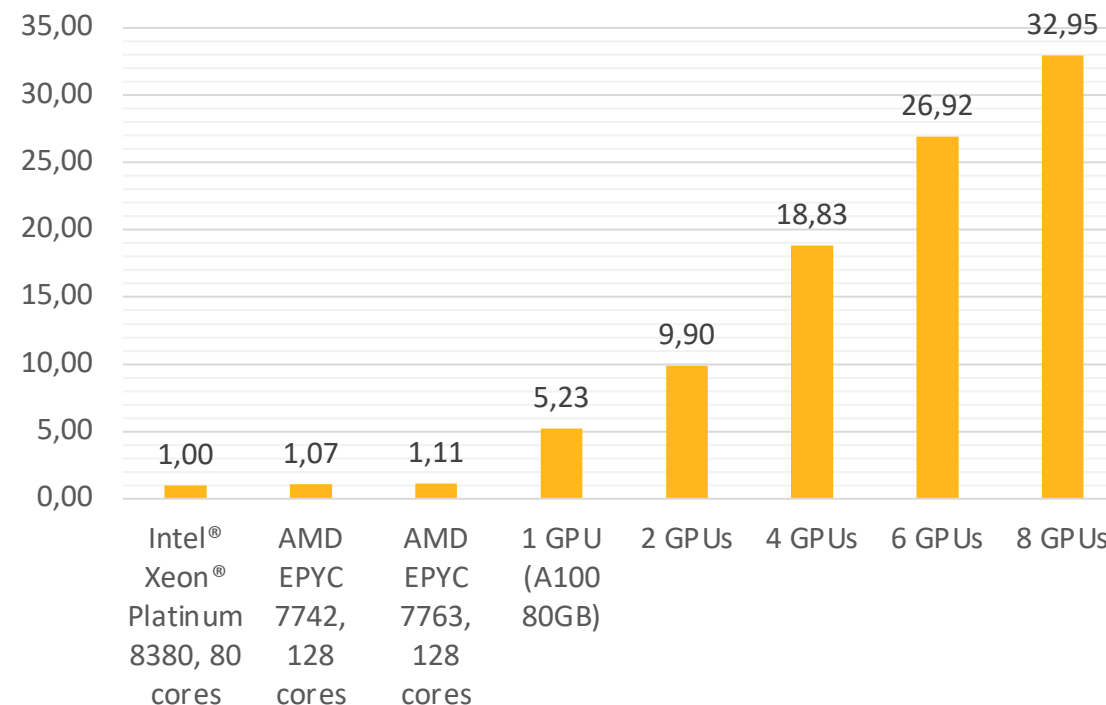

Mechanical



HFSS/SI-Wave

Native GPU Implementation Shows Astounding Performance Gains

- 105M case single precision with GEKO runs on 1 A100 80GB card
 - ~1 iteration per 2 seconds, 20 minutes to converge such a case from initialization
- 1 A100 GPU ≈ 500 Intel Icelake cores
- 8 A100 GPUs ≈ 3000 Intel Icelake cores
- Parallel efficiency is 80% from 1 to 8 GPUs
- Compared hardware, run by Nvidia performance lab
 - Platinum 8380@2.3GHz 3.4GHz Turbo (Ice Lake)
 - AMD EPYC 7742@2.25GHz 3.4GHz Turbo (Rome)
 - AMD EPYC 7763@2.45GHz 3.525GHz Turbo (Milan)
 - NVIDIA DGX 8 x A100-SXM4-80GB(GA100)



“One DGX will replace 30 high-end dual CPU servers”


One DGX will replace 30 high-end dual CPU servers




- / 4 typical GPUs > 1,000+ CPU cores
- / 7x cheaper hardware purchase cost
- / 4x lower power consumption

Ansys / FLUENT




Live-GX SOLVER




 **NVIDIA.** Early access to the **Hopper Architecture (H100)**: 24 more SMs and much higher mem bandwidth than A100

DrivAER Car - CPU/GPU – LES Smag. – Surface Cp values


- Rapid octree mesh, 19.43M cells, single prism layer
- Good agreement with expt. and CPU solution over probes 1-35
- Reasonable agreement elsewhere
- Drag overpredicted on this mesh with GPU solver
Iso. Vort. Mag = 500 [s⁻¹]




Generic Combustor: Strong Scaling



24M Cells
MCIUPS (Million Cell Iteration Updates per second)



115M Cells
MCIUPS (Million Cell Iteration Updates per second)



- ~75% parallel efficiency with 8 GPUs
- Runs ~1 iteration/second on 1 GPU

- ~83% parallel efficiency with 8 GPUs
- Runs ~1 iteration/second on 6 GPUs

Single-GPU Performance Across Various Hardware Generations


- Fluent Intel 48-core Sky Lake system is used as the basis
- Tesla P100 is about 2.5X as fast
- Tesla V100 is about 4.5X as fast
- Tesla A100 is close to 8X as fast




nodes: 2,352,640
 faces: 6,663,296
 cells: 2,155,529

Car 2M Hex-poly mesh, Flow + Turbulence, single GPU


Devices	Compute	Memory (GB)	Memory bandwidth (GB/s)	Cores
Fluent CPU 48 Skylake cores		192	107.5	48
Quadro P6000	6.1	23.9	413.2	3840
Tesla P100	6	15.9	698.2	3584
Quadro GV100	7	31.7	830.1	5120
Tesla V100	7	15.8	856.4	5120
Tesla A100	8	39.6	1483.2	6912



Strong Scaling with 25M Car Case, Poly-Hexcore Mosaic™ Mesh



DrivAer 25M, 2 x Intel(R) Xeon(R) Gold 6242
CPU
vs A100



Ansys SimAI: Artificial Intelligence for Simulation

Ansys AI

Ansys AI+

AI Add-ons to Ansys products across portfolio




AI Add-ons to various Ansys simulation products that enhance simulation

Various Improvements

Ansys SIMAI

ML platform for simulation across the physics



Extremely fast and reliable physics predictions which learns from existing data

10x to 1000x Faster

Ansys GPT

Virtual assistant to Ansys products




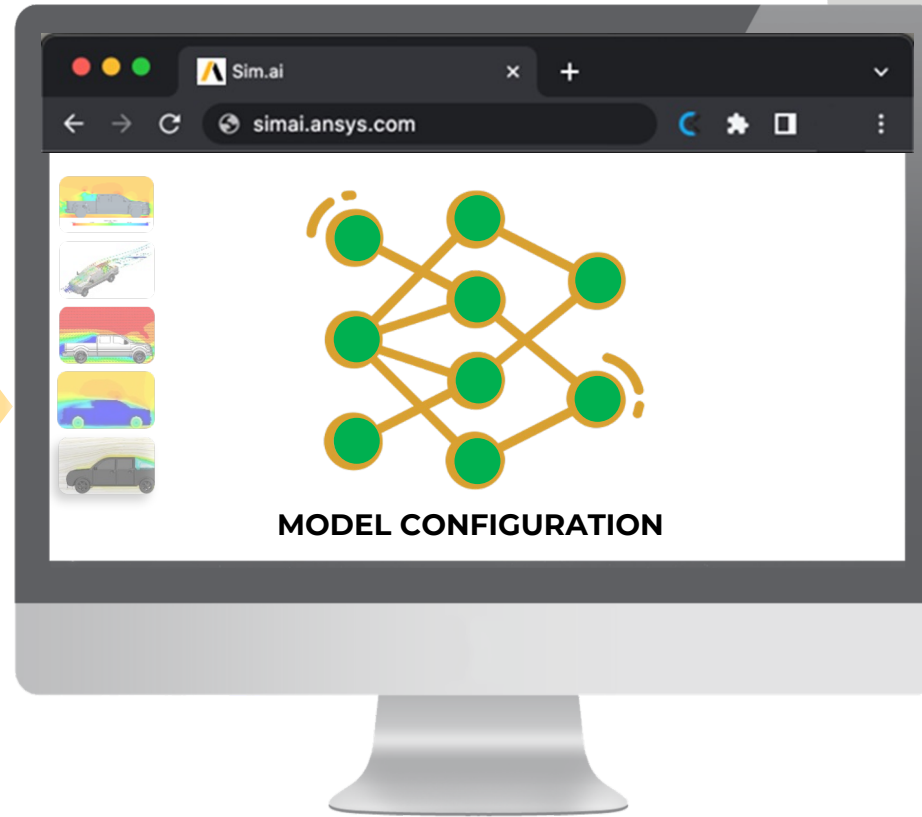
Natural language assistant for documentation, training

Simple & Natural UX

Predict at the Speed of AI

simtec

ANSYS SIMAI



1- UPLOAD
Your Past Data

2- TRAIN
Your AI Model

3- PREDICT
In Seconds

ANSYS

Unleash the Power of AI for your Design with Ansys SimAI

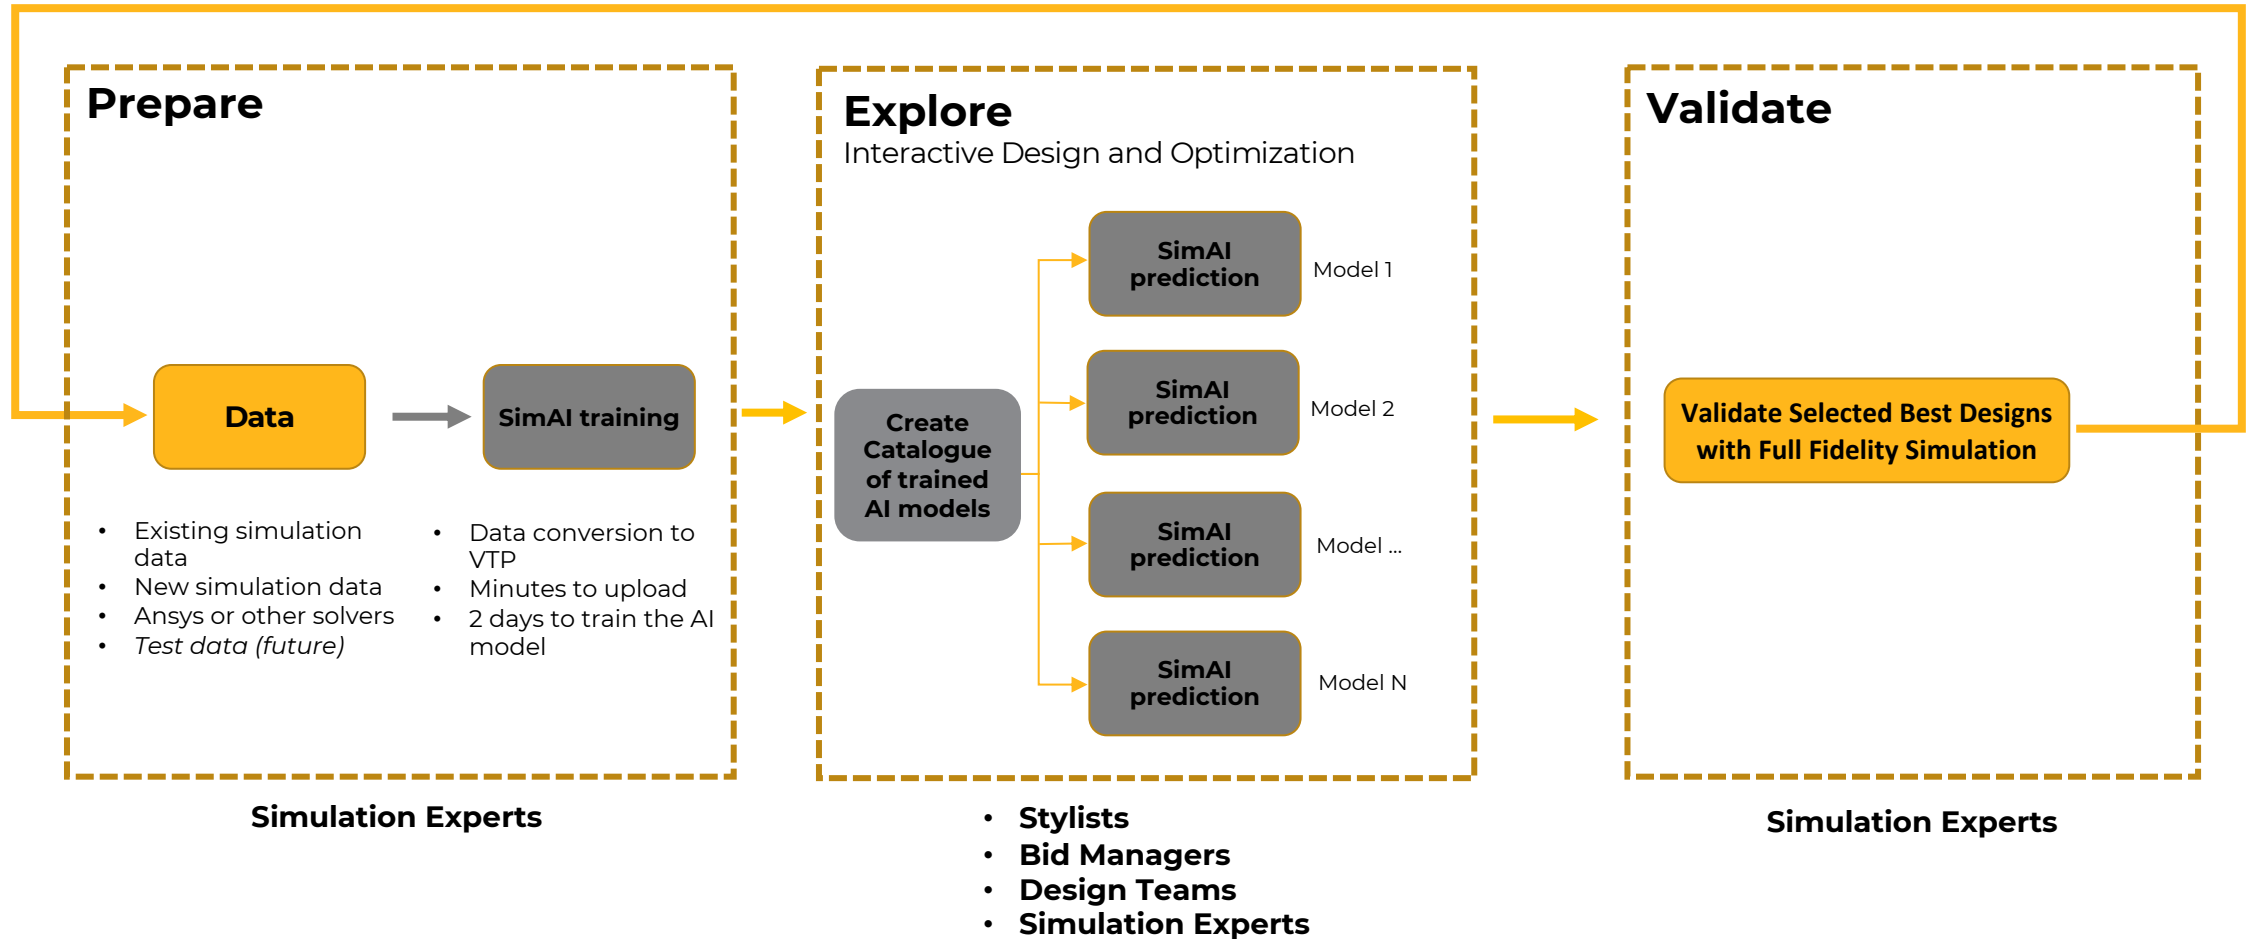
Ansys SimAI is our new cloud-based AI platform for Simulation:

- Train the AI without having to parametrize your geometry
- Predict performance across design changes, even when the geometry structure is inconsistent
- **Leverage previously generated simulation results to train the model**

Ansys SimAI is physics-neutral:


- Any physics – Fluids, Structures, Emag, Optics
- Across all industry segments – Aerospace, Automotive, Semiconductor, etc...
- Works with any 3D simulation data, whether it is Ansys or not






Ansys SimAI – In your Organization

3D simulation and test results



Your catalog of trained AI models



AI prediction



Upload your data

1min per Gb

Train your AI models

48h


Predict new designs

30s

Fluids



Fluent




CFD comparison
Thermal management
Cooling design

Structures



Mechanical
LS-DYNA




Generative design
Wire forming process
Impact performance
Crystal plasticity homogenization
Stress + deformation

Electronics



HFSS
Maxwell
Icepak

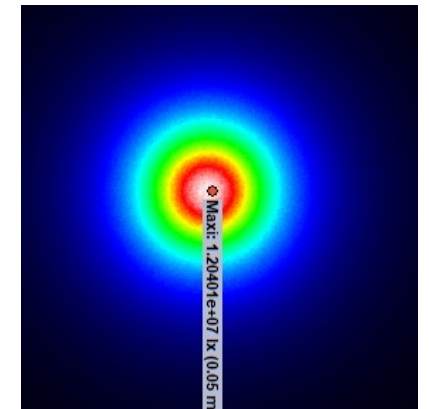


Antenna design & placement
Magnet placement
PCB EM losses and forces
Electric motor design

Optics



Speos



Illumination

SUV Aerodynamic Performance



Challenge

CO2 emission reduction plans (e. g. [WLTP](#)) require automotive manufacturers to **assess the aerodynamic performance of all design variants** of a new car family. Wind tunnel tests are too slow and costly to solve this challenge. And it is important to **shrink the development time of the next generation of electric cars**.



- Drastically faster predictions of aerodynamic performance are required to stay competitive.

Solution

- ~ 50 accurate CFD results are used to create the AI model, including car exterior shape variations and topological changes (rear mirror, ski rack, spoiler, etc). (*Fluent, SimAI*)
- **SimAI Prediction on new SUV geometry in less than 1 min.** (*SimAI*)
- SimAI Drag error compared to CFD: **less than 0.5% (5 to 10 drag counts)** and accurate skin friction field and wake topology prediction. (*SimAI*)

Benefits

- Assess **more car designs: 20x** compared to traditional simulation methods and optimize quicker.
- **Predict consistent aero performance across design changes faster (10 to 100x)**, even when the geometry structure is inconsistent by leveraging on past CFD simulations database (earlier design phase, previous car generations).
- **Shift Left:** Cut-down your design process duration and cost by allowing **designers to use fast and meaningful aero prediction**.



Predicting Traction Motor Performance for Electric Powertrain

Challenge

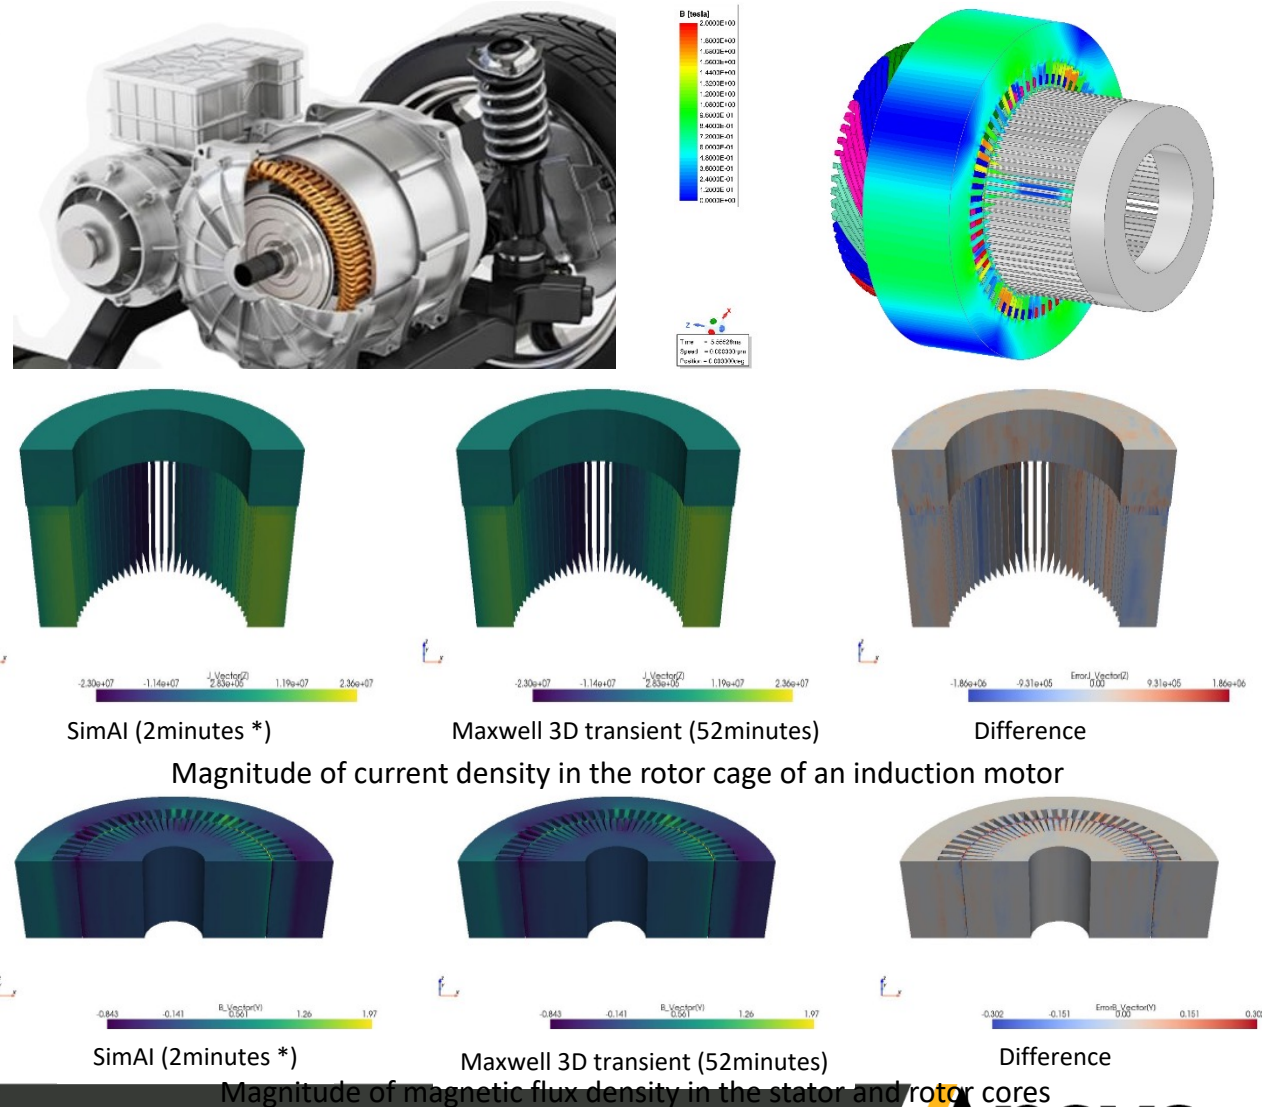
- Conduct rapid prediction of motor EM torque, efficiency, losses, and forces/stress with design changes.
- Reuse historical data from previous production design to derive insights for next generation design at early design stage.
- Explore alternative materials, innovate more compact motors.

Solution

- Training dataset: **tens to hundreds of accurate simulation results** are used to create AI models of different fidelities. (*Maxwell* and *SimAI*)
- **Each new design is evaluated in minutes with the AI model build.**
- **Smooth magnetic field and electric current distributions** are accurately predicted by the AI model.
- Good accuracy of global coefficient for total loss > 95%.

Benefits

- **Faster Time to Market:** Quick design space exploration early in the design cycle of the motor: 20X to 50X faster to conventional simulation techniques.
- Obtain **more innovative designs** by leveraging historical insights.




/ Testimonial – Renault Group



With Ansys SimAI, we will be able to easily test a design within minutes and rapidly analyze the results, **ultimately redefining our digital engineering workflow** and reshaping our perception of what is possible. By enhancing simulation speed, **we can explore more technical possibilities** during the upstream phase of our projects and reduce the overall time-to-market.”

William Becamel

Expert Leader in Numerical Modelling and Simulation | Renault Group



Source: <https://www.ansys.com/news-center/press-releases/1-9-24-ansys-launches-simai>

SLING



EuroHPC
Joint Undertaking



REPUBLIKA SLOVENIJA
**MINISTRSTVO ZA VISOKO ŠOLSTVO,
ZNANOST IN INOVACIJE**

Projekt EuroCC 2 financira Evropska unija. Financiran je s sredstvi Skupnega podjetja za evropsko visokozmogljivo računalništvo (EuroHPC JU) ter Nemčije, Bolgarije, Avstrije, Hrvaške, Cipra, Češke republike, Danske, Estonije, Finske, Grčije, Madžarske, Irske, Italije, Litve, Latvije, Poljske, Portugalske, Romunije, Slovenije, Španije, Švedske, Francije, Nizozemske, Belgije, Luksemburga, Slovaške, Norveške, Turčije, Republike Severne Makedonije, Islandije, Črne gore in Srbije v okviru sporazuma o dodelitvi sredstev št. 101101903.