Using HPC for ANSYS

Klas Johansson

Agenda

- The ANSYS family 2020
  - A short tour of some of the latest additions (Discovery, optiSLang, LSDyna, Twinbuild, Ensight, Granta)
- HPC Trends in the industry
- HPC Benchmarks of some ANSYS tools

ANSYS Discovery

The Design Challenge

Product development cycles are getting shorter while product complexity continues to increase. The company which can explore and refine the most design ideas in the shortest time wins.

Ideation

Digital exploration of many diverse options to uncover promising ideas and weed out unpromising ones

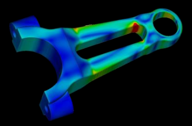

Refinement

Detailed design to narrow down specifics & optimize performance

Validation

High accuracy analysis to certify and test systems

EDR™ MEDESO

The Current State of Simulation

The majority of simulations are performed late in the design cycle by a small group of expert analysts.

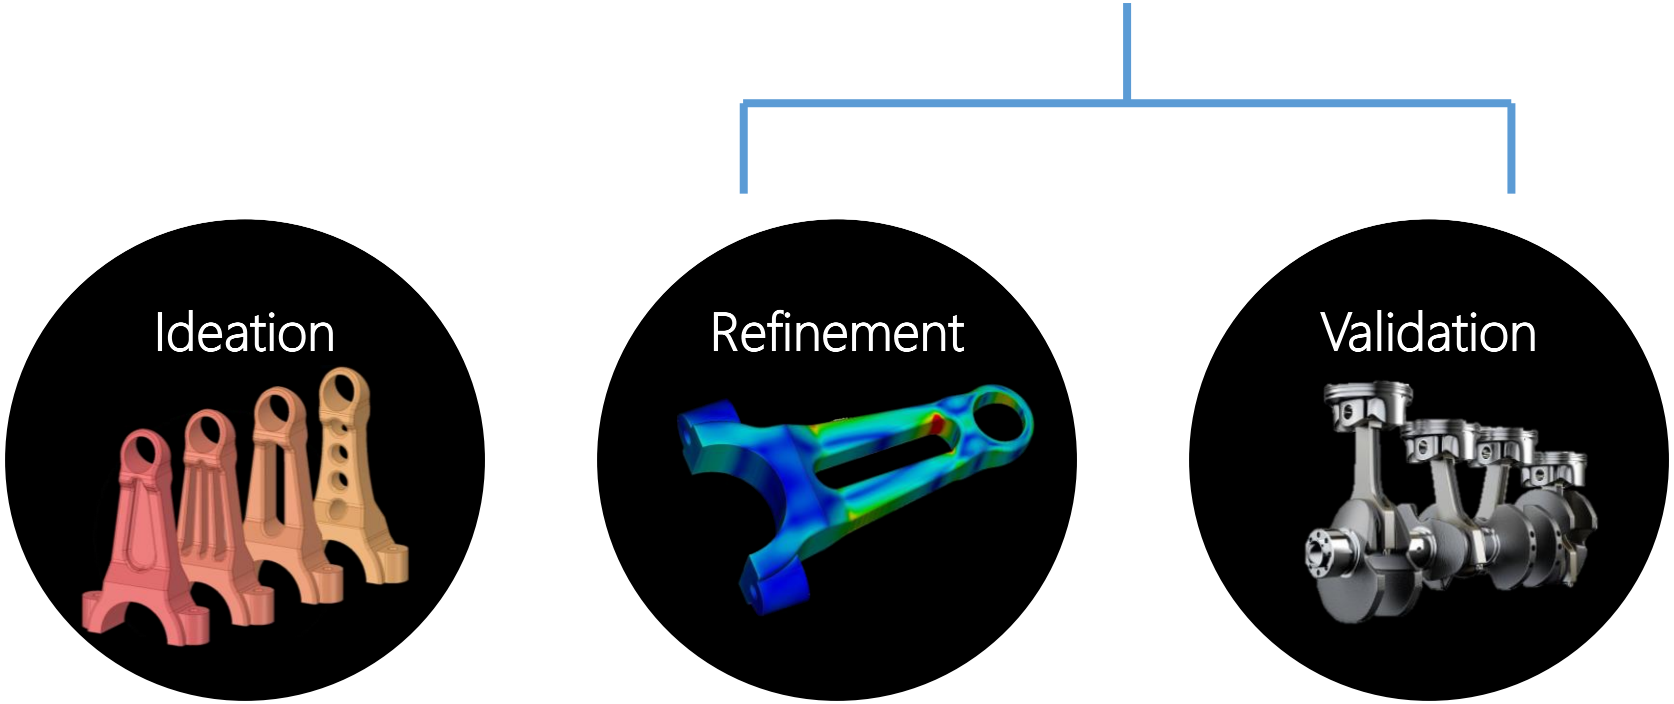

ANSYS Discovery

Designed for extreme ease of use

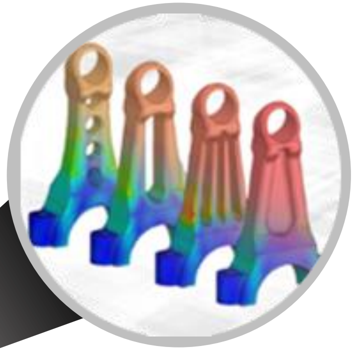

Instantaneous simulation for rapid design exploration

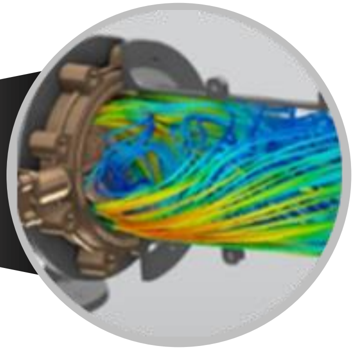

High fidelity simulation for detailed analysis of product performance

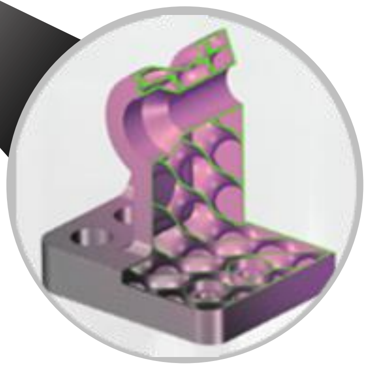

Multiple physics and rapid concept modeling in a single platform

EDR MEDESO

ANSYS Discovery in Action

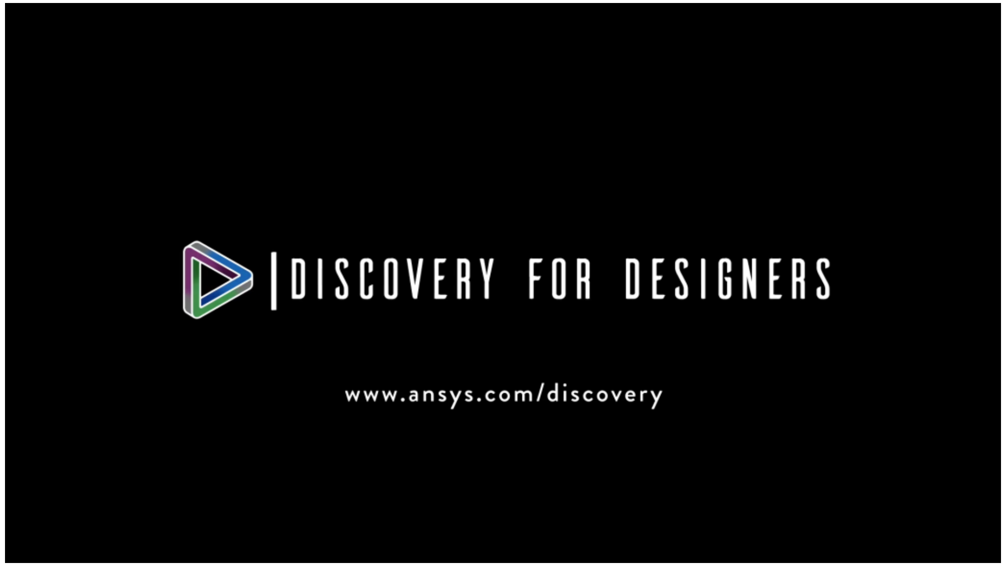

Overview

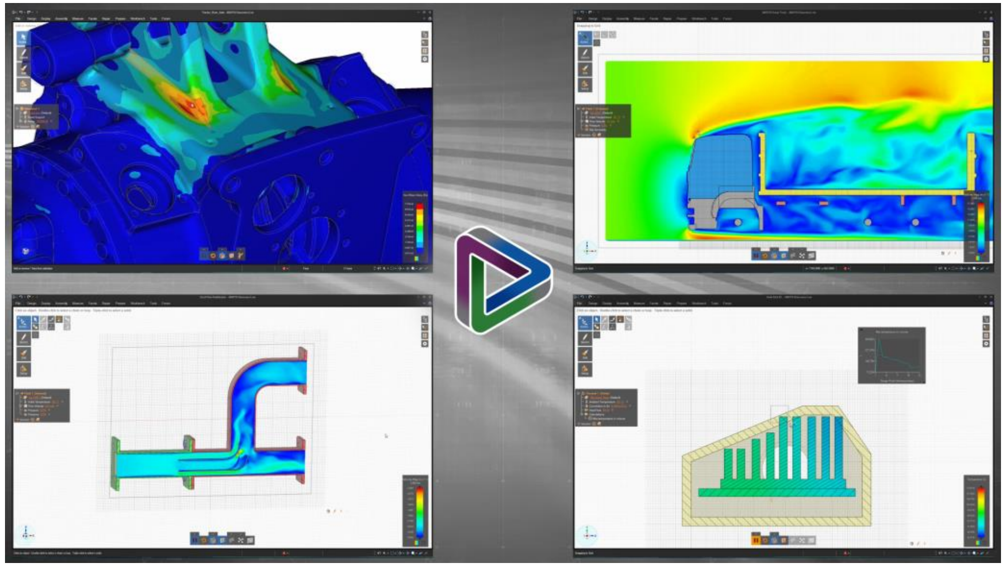

Demo

EDR™ MEDESO

ANSYS optiSLang

EDR™
MEDESO

ANSYS optiSLang

- General purpose tool for variation analysis
- CAE-based design points (and/or data points) for the purpose of
  - sensitivity analysis
  - calibration of virtual models to physical tests
  - design/data exploration
  - optimization of product performance
  - quantification of product robustness and product reliability
  - **Robust Design Optimization (RDO)** and Design for Six Sigma (DFSS)

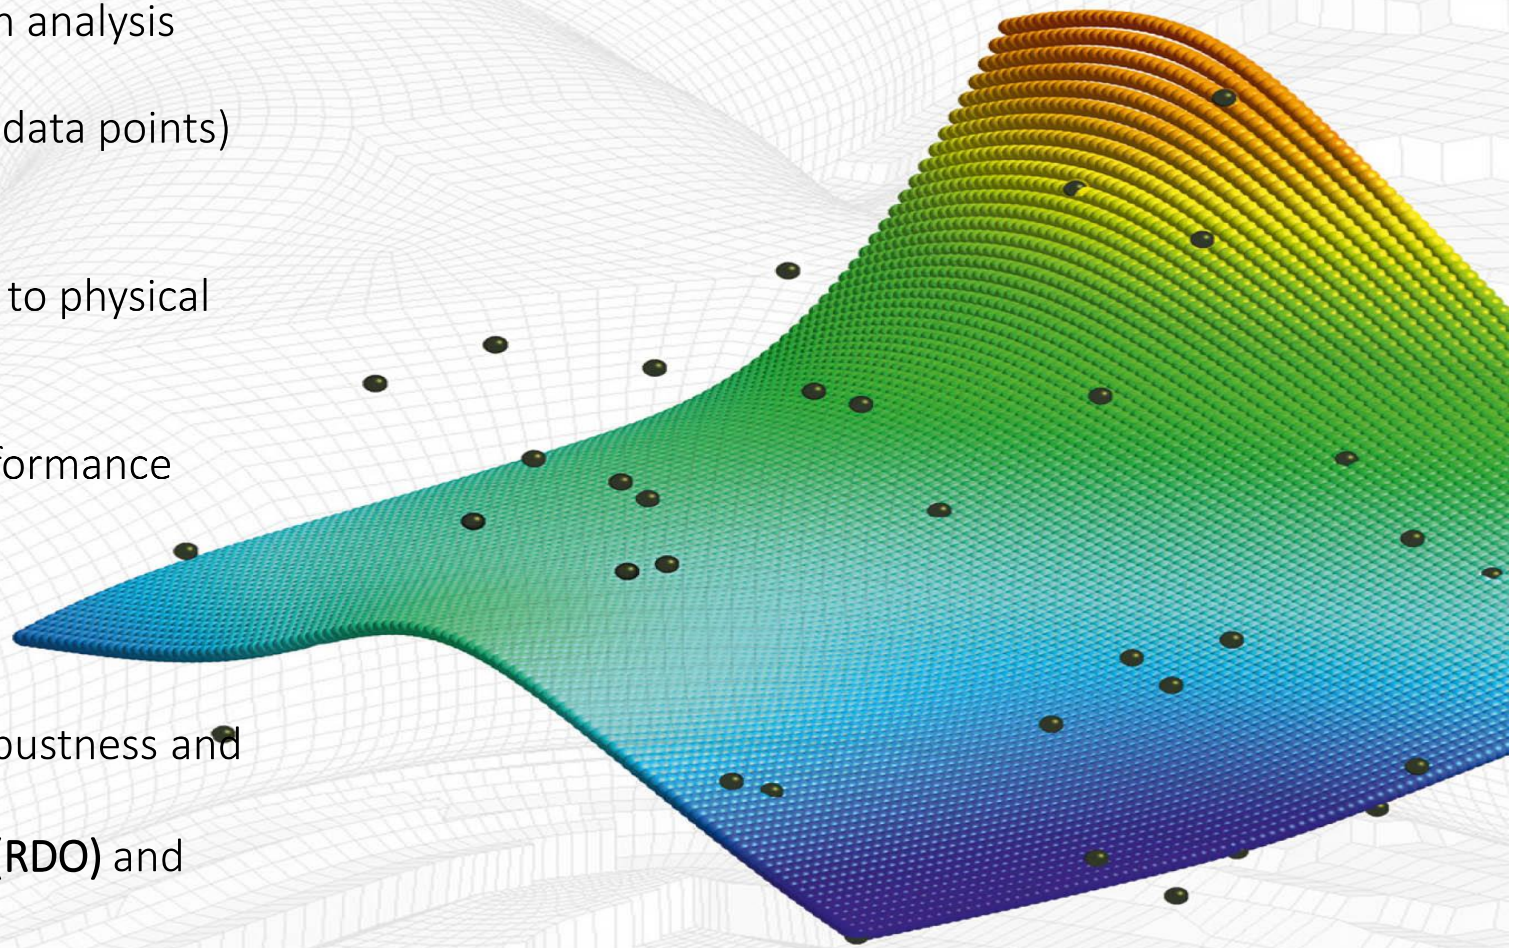

EDR™ MEDESO

ANSYS optiSLang

- Serves all ANSYS tools and HPC/Cloud components!
- Available within ANSYS Workbench in the form of a well integrated add-in at the project page
  - Extremely easy to use!
- Standalone GUI mode which supports process integration, automation and workflow generation of ANSYS and any 3rd party code

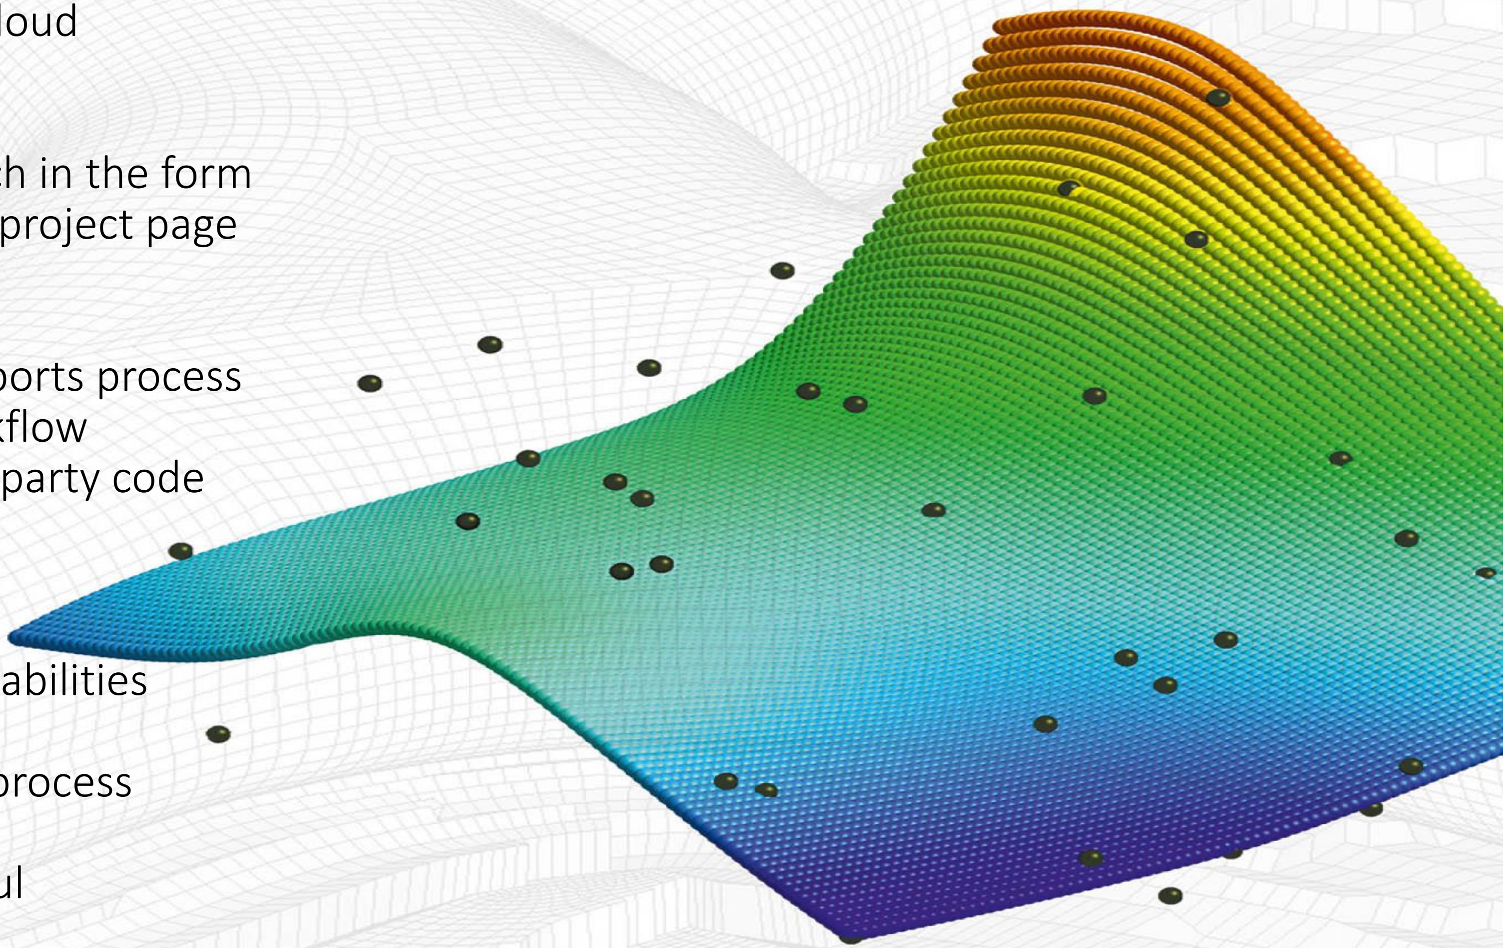

- Advanced process building capabilities
- Reserve the ownership of the process
- optiSLang & SPDM is a powerful combination!

EDR™ MEDESO

ANSYS optiSLang's philosophy

- Minimal user input
- Automatic reduction of parameters
- Robust and efficient, even with noisy data or design point failures
- Built-in mechanisms to ensure robust workflow
  - Continue crashed/interrupted session
  - Recalculate failed designs
  - Handles up to 50% design point failure rate

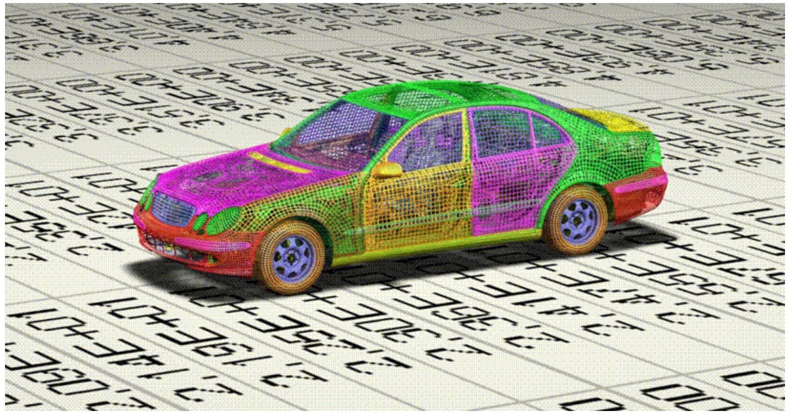

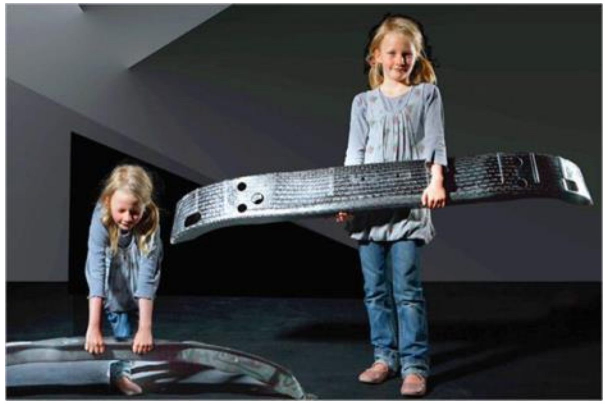

EDR™ MEDESO

ANSYS optiSLang - ease of use

- Engineers and Designers should not have to choose from a list of detailed settings and complicated algorithms!
- ANSYS optiSLang's functionality is compressed to three wizards (sensitivity, optimization, robustness).

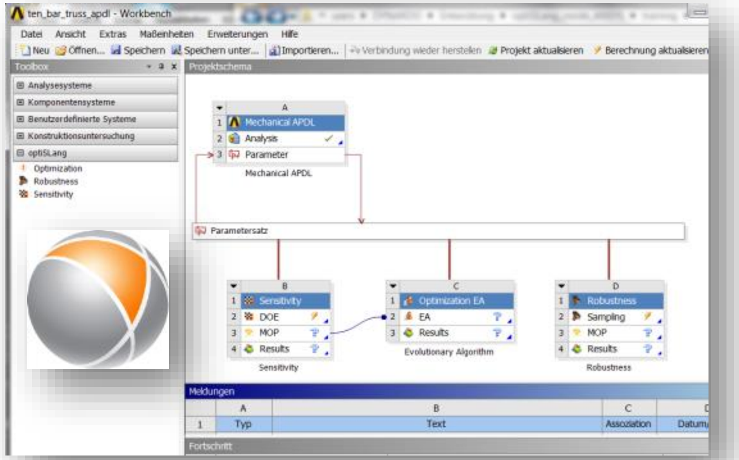

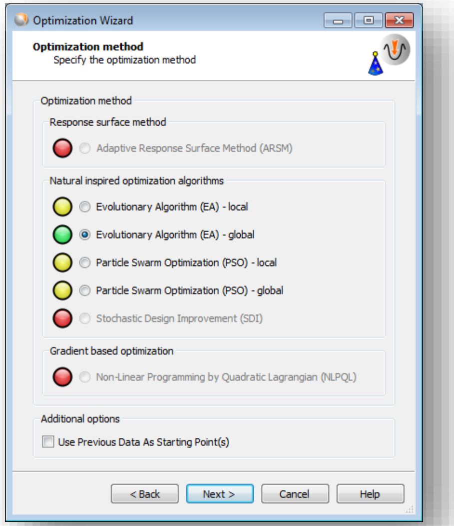

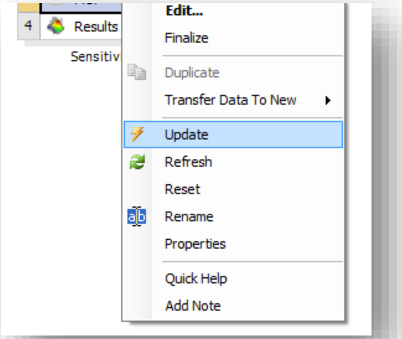

EDR MEDESO

ANSYS optiSLang - MOP

MOP – Metamodel of Optimal Prognosis

Solves three important tasks:

- Determination of relevant parameter subspace
- Determination of optimal approximation model (RS)
- Present a single value for response surface quality (coefficient of prognosis - CoP)

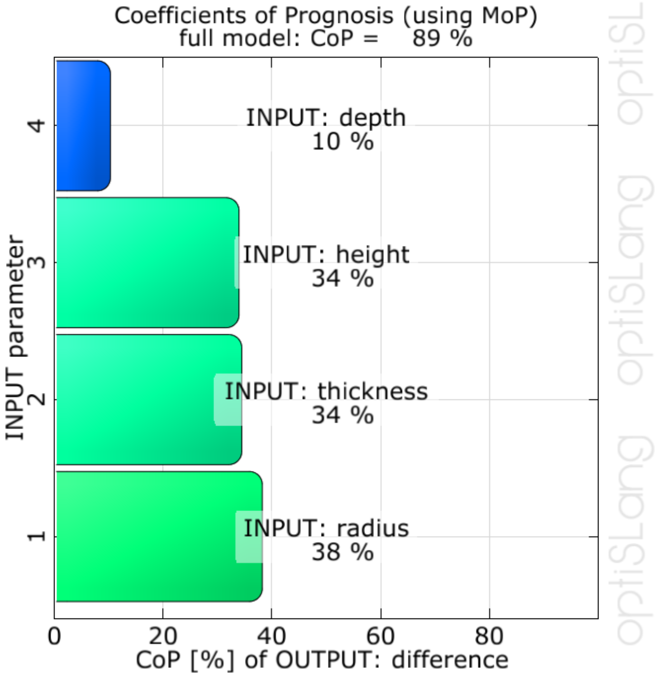

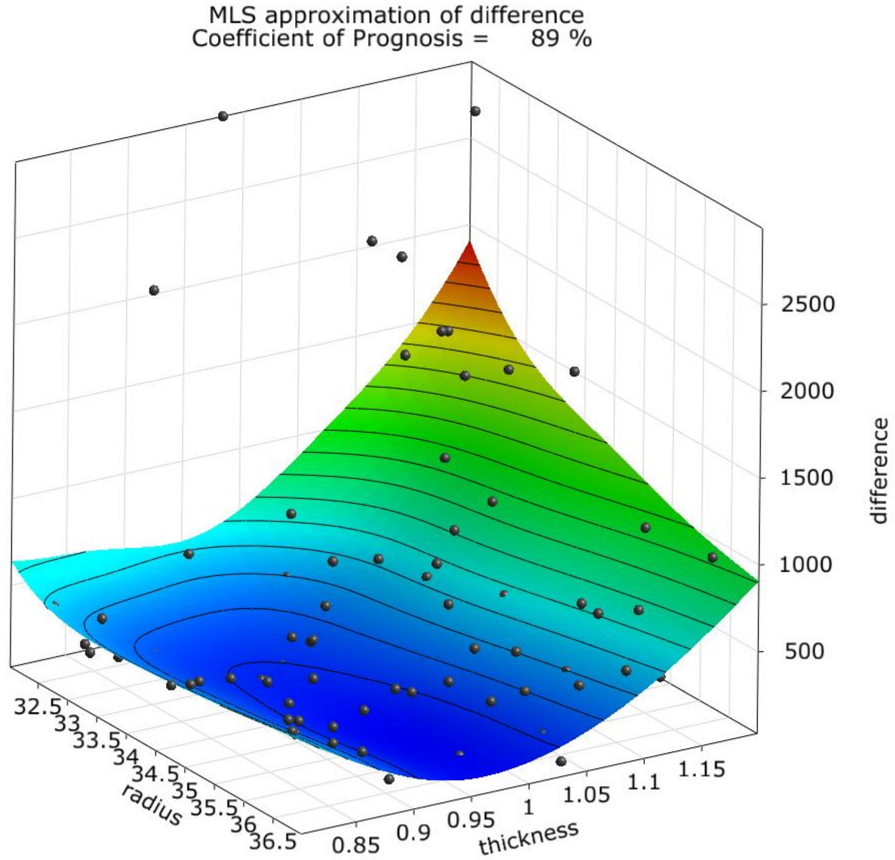

ANSYS LS-dyna

What is LS Dyna?

- LS-Dyna is the most powerful and flexible general purpose explicit analysis code in the world.

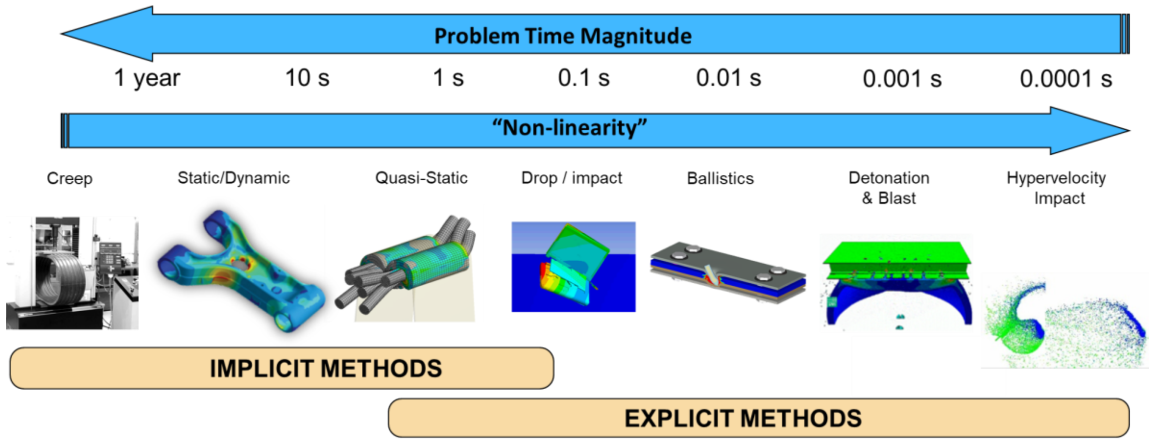

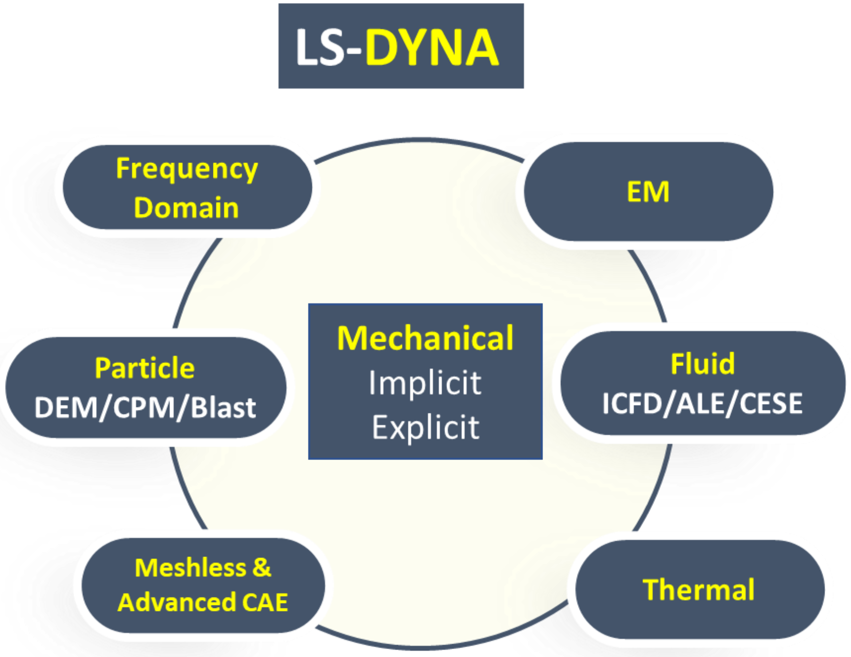

- Its core explicit solver can be loaded with forces generated from several other supporting solvers which include CFD, thermal, electromagnetics and particles.

EDR™ MEDESO

What Are The Strengths Of ANSYS LS-Dyna?

- Very fast, very efficient parallelization up to hundreds of processors
- More than 200 nonlinear material models
- Element types can be chosen for fast solution time or high accuracy
- Joints, bolt pretension, contact interference
- Keyword snippets: Solver has many more features than Workbench Interface exposes today
- Thermal loads
- ANSYS Mechanical integration provides unparalleled ease of use, parameterization, simple workflow and meshing all in one environment.

Automotive Crash & Occupant Safety NVH & Durability Battery Reliability	**Civil Engineering** Blast Proofing Earthquake Safety Tents
Aerospace Bird Strike Blade-off Containment Crash	**Electronics/Hi-Tech** Drop Analysis Package
Manufacturing Forming Stamping Machining	**Defense** Projectiles and Weapons Blast and Penetration Underwater Shock Analysis
Consumer Products Packaging Switches	**Bio-Medical** Devices & Equipment Medical procedures

ANSYS Twinbuilder

ANSYS Simulation Platform Overview

… To Complete Systems Simulation

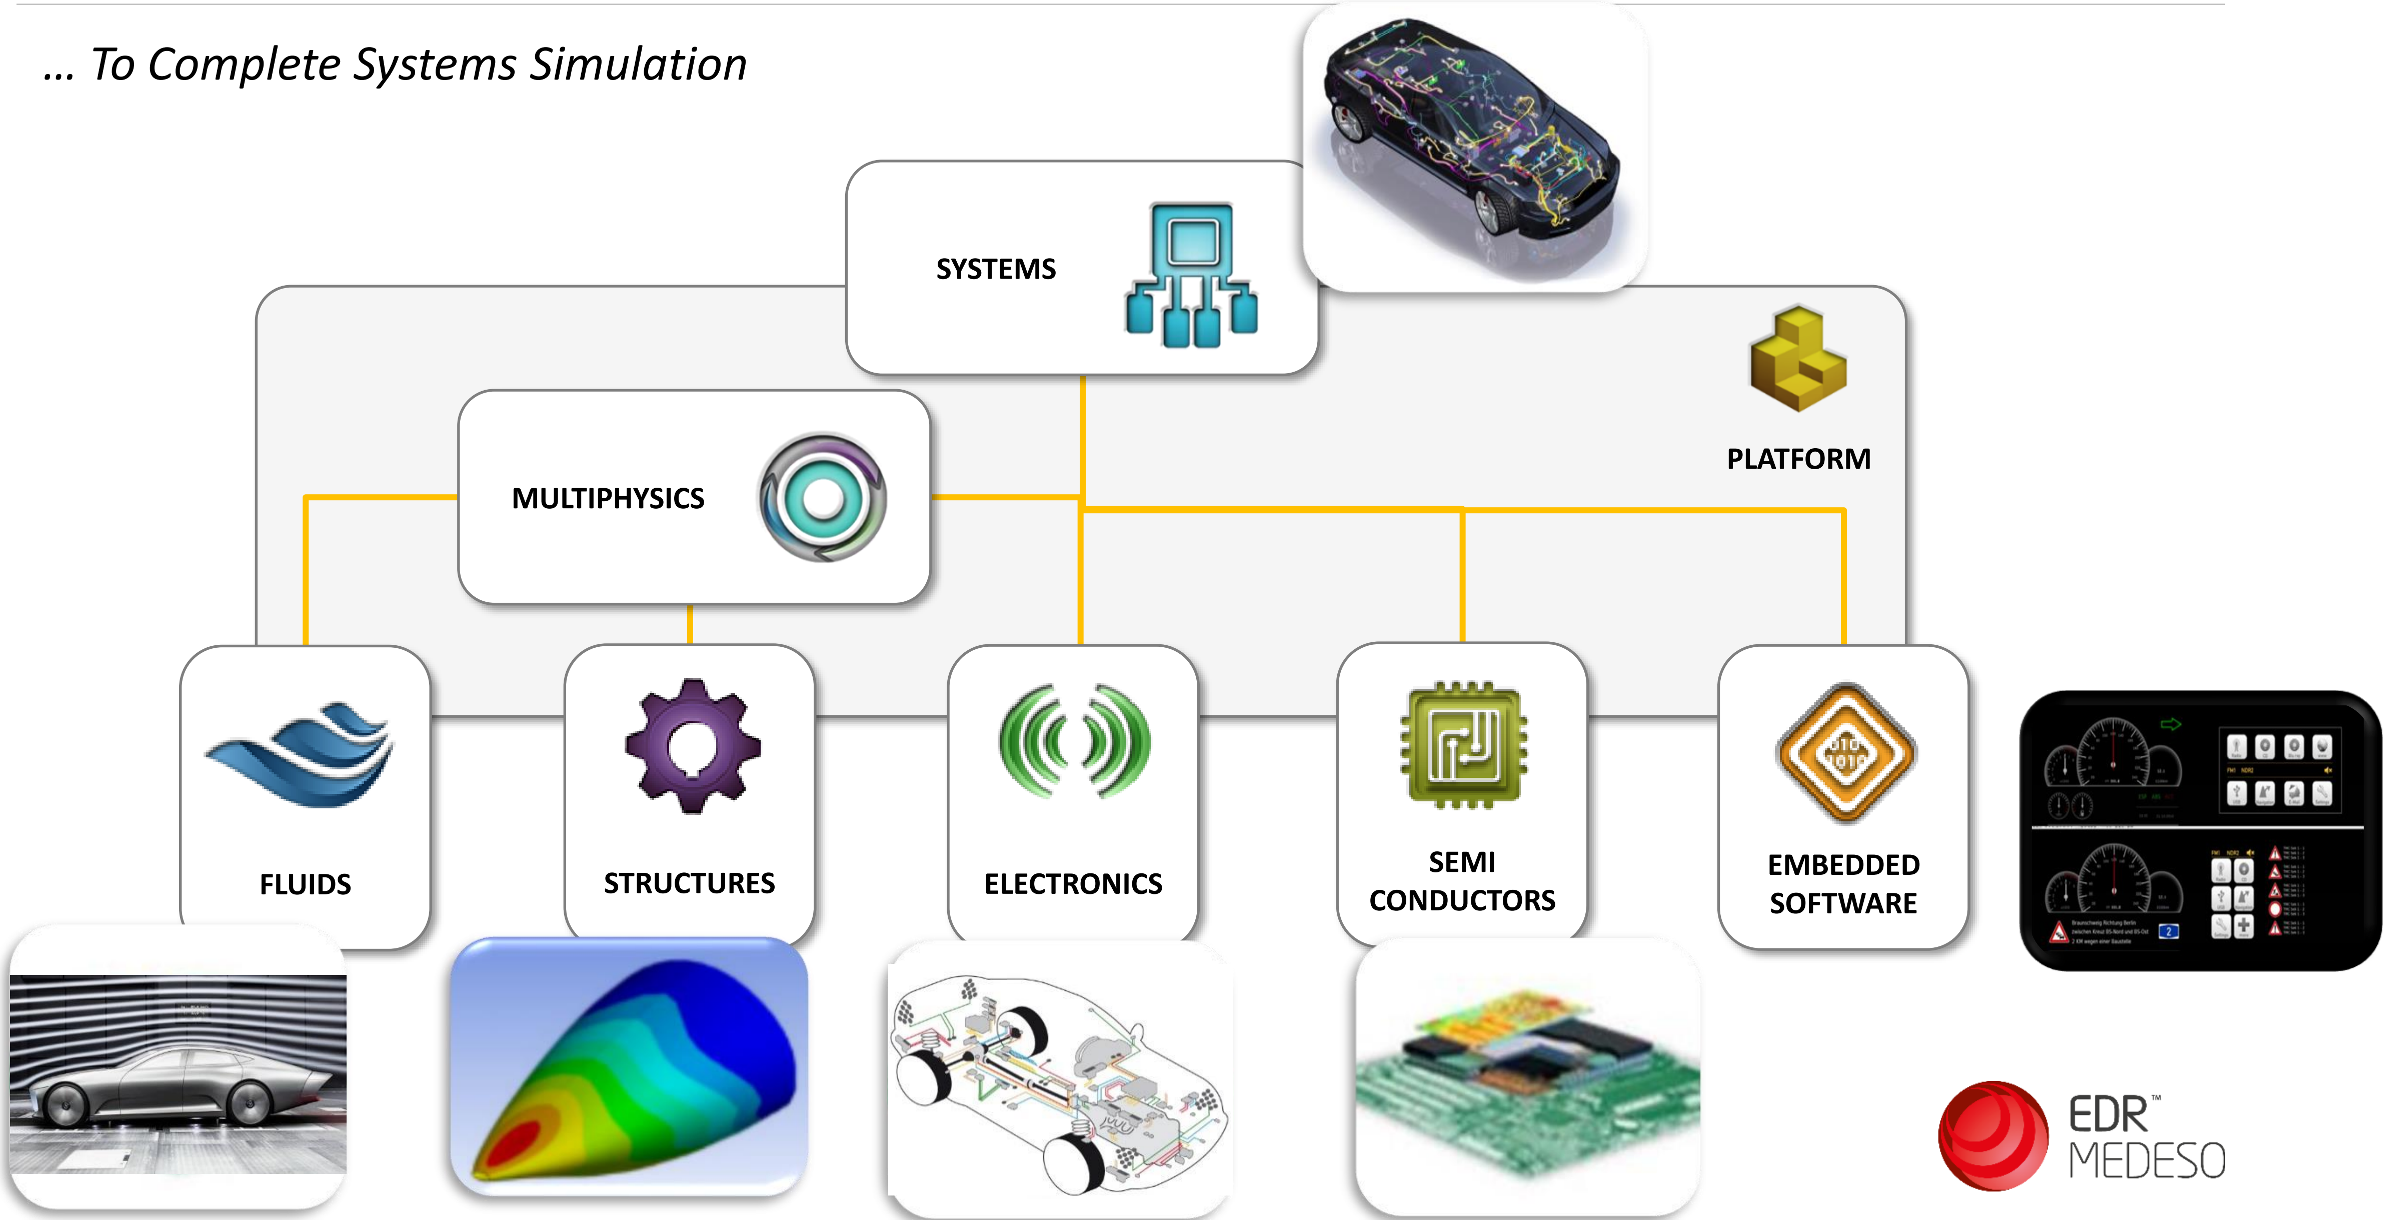

EDR™ MEDESO

Key Capabilities of ANSYS Twin Builder

Build

Build an accurate Physics-based Digital Twin in record time

System Simulation

Validate

Validate and Optimize the Twin

System Validation and Optimization

Deploy

Connect the Twins to IIoT Platforms and Deploy Run times in operation

System Predictive Maintenance

IIoT Platform connectivity

- Quickly connect to popular IIoT platforms
- Export from Twin Builder to generate portable, cloud deployable Twin

IIoT Platforms

ptc

PREDIX

SAP

Connect

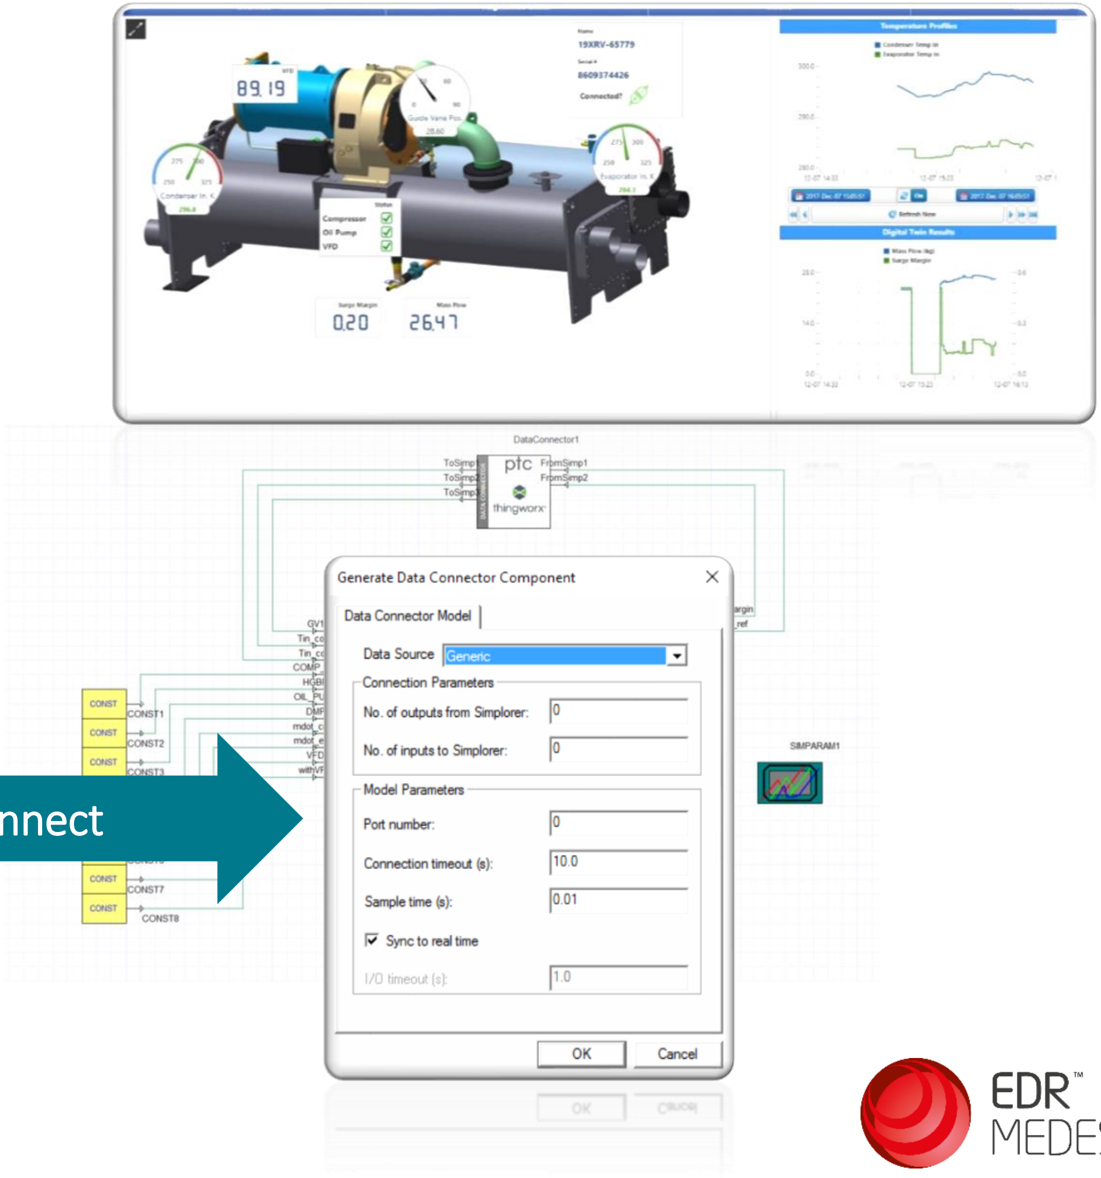

ANSYS Ensight

ANSYS Ensight

- A general post processing package aimed at providing a single consistent environment to post process any solver, any physics, any solution.

- Geared towards large model post processing, scripting, automation of processes.
- Long list of import data formats & solver data reading.
- Advanced rendering features (Volume Rendering, LIC, Ray Traced Scenes, Keyframe Animation)
- Multi-Case interactive comparison.

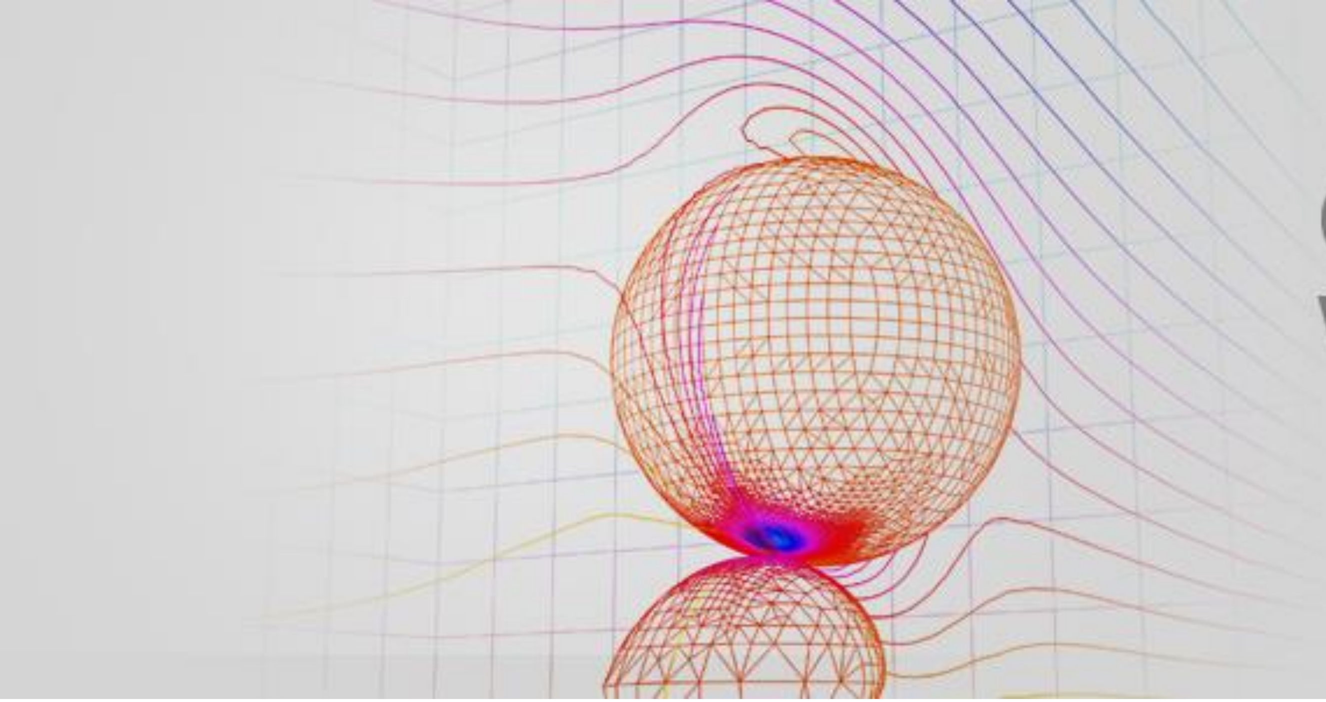

Can EnSight work with your data?

Yes. It can.

ANSYS
- CFX
- Forte
- Fluent
- Icepak
- Polyflow
- FENSAP-ICE
- Mechanical
- AUTODYN
- LS-Dyna

CFD
- AcuSolve
- CFD++
- CGNS
- Cobalt
- CONVERGE CFD
- FAST
- FIDAP
- FIRE
- Flow-3D
- GASP/GUST
- KIVA
- NASTAR
- OpenFOAM
- Overflow
- PAM-FLOW
- Plot3D
- PowerFLOW
- RADIOSS-CFD
- SC/TETRA
- Star-CD, CCM+
- VECTIS

FEA
- ABAQUS
- I-DEAS
- LS-DYNA
- MP-Salsa
- MSC.Dytran
- MSC.Nastran
- MSC.Marc
- MSC.PATRAN
- NX Nastran
- PERMAS BIF/BOF
- RADIOSS

Other
- CTH
- Excel/Flatfile
- HDF5
- EXODUS/PXI
- SILO
- MESHTV
- Netcdf
- MFIX
- MRI data

DEM
- Barracuda
- EDEM
- Particle Works
- SPH
- LAMMPS
- LIGGHTS

EDR™ MEDESO

EnSight Standard & Enterprise

- EnSight works on a client-server architecture that distributes the workloads between the client and one or more servers – this architecture leaves the data where it was computed

Client (GUI & all graphics)

Server (cluster): all data, extraction algorithms; local or remote, SMP or DMP

- The GUI and graphics are handled by the Client; the Server takes care of all data and the data extraction algorithms; the Server can be local or remote

EDR™ MEDESO

ANSYS Granta

Materials: the next step for digital transformation

When you make a product, you consider four factors

	Shape	Function	Manufacturing	Material
Industrial era		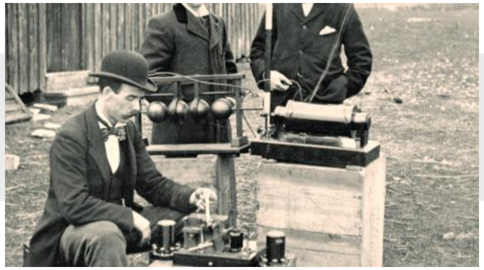	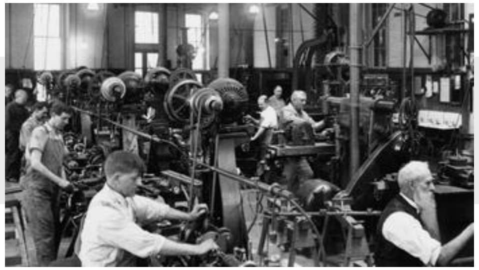	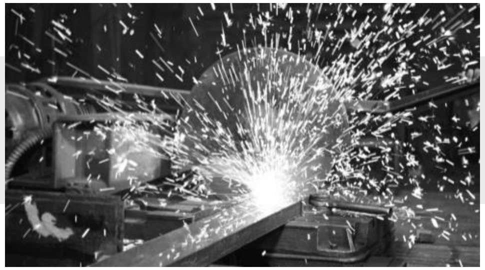
Digital transformation	 CAD	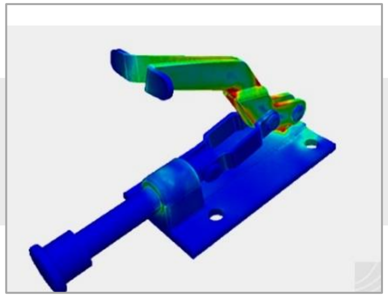 CAE	 CAM	GRANTA MATERIAL INTELLIGENCE Materials information technology

EDR™ MEDESO

ANSYS Granta

Strong pedigree

Cambridge University spin-out,
R&D record, passion for materials…

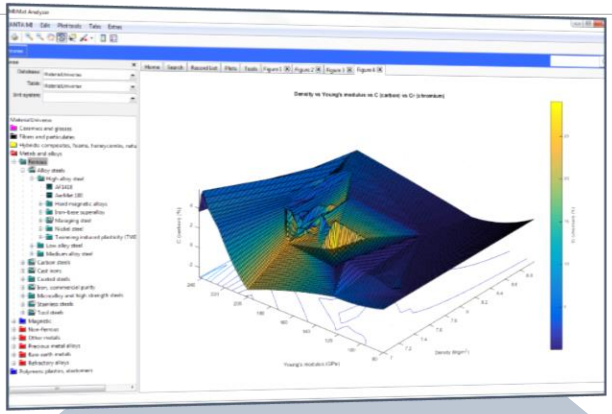

Materials Information Technology

Exceptional collaboration

Customers, Industrial Consortia, Educational Symposia, strategic partnerships`

Software	Information	Services	Network

EDR™ MEDESO

The different tools

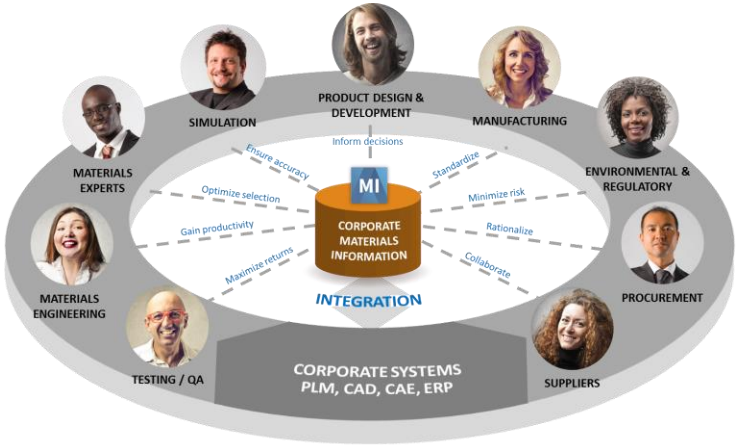

GRANTA MI™

The leading system for materials information management in engineering enterprises

The platform for material intelligence

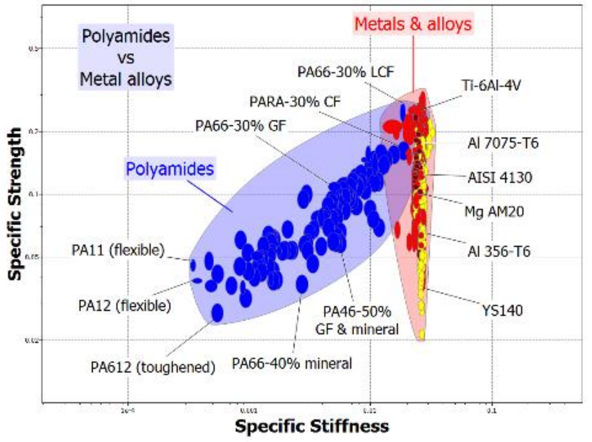

GRANTA Selector™

Smart materials decisions

PC product for materials experts

GRANTA EduPack™

World-leading resources for materials teaching in engineering, science, and design

EDR™ MEDESO

HPC Trends and development

Globalization and HPC Scale-up are Driving a Paradigm Shift in the Deployment of Simulation

HPC infrastructure and **data** is evolving to **centralized** and **consolidated resources**, with users **remote, global,** and **mobile**

Driven by:

- Scaled up **demand** for **high performance computing (HPC)** to support **higher fidelity** simulation results
- **Operational efficiency** and **asset optimization**
- Need for **collaboration** and **data sharing** across geographically **distributed users**
- Increasing focus on simulation **IP protection** and **process traceability**

EDR™ MEDESO

Leading Companies are Consolidating Simulation in the Data Center, with Users Remote and Mobile

Some examples:

GM… "goal is to make data easier to access and use, and to integrate it better across groups and divisions to reveal more insights."

Cummins… "we can improve our customers' business through real-time optimization of the powertrain utilizing off-board computational resources."

RCR… "being able to get more cloud compute capacity is exactly what we're after. As our mesh grows, we don't have to invest in increasing our on-site cloud compute capacity."


Automotive News

HOME CARS & CONCEPTS NEWS BY BRAND DEALERS OEM/SUPPLIERS MOBILITY SPECIAL REPORTS OPINION VIDEO DATA WEBINARS EVENTS JOBS

You are here: Home » Technology » General Motors » GM It Transformation »

FIRST OF A 2-PART SERIES

GM plots next phase of IT overhaul

Unlocking the potential of a vast data empire

September 18, 2017 @ 12:01 am

Michael Wayland

1 Shares

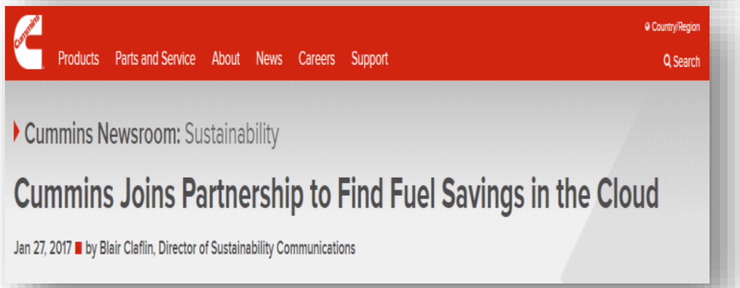
Country/Region

Products Parts and Service About News Careers Support

Search

Cummins Newsroom: Sustainability

Cummins Joins Partnership to Find Fuel Savings in the Cloud

Jan 27, 2017 by Blair Claflin, Director of Sustainability Communications


Richard Childress Racing Leverages ANSYS to Improve Racecar Speeds

ANSYS

Multi-year partnership agreement fast-tracks RCR's aerodynamics and engineering programs

NEWS PROVIDED BY

ANSYS, Inc. →

Feb 15, 2018, 07:00 EST

SHARE THIS ARTICLE

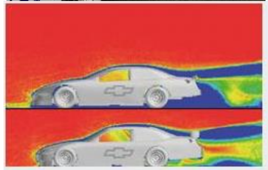

EDR™ MEDESO

Implementation of Distributed Simulation Requires a Focus on End-user Process and Data

Best Practices

Consolidate HPC and data

- Operational efficiency
- Avoid moving data
- Encourage collaboration

Enable effective remote access

- Beyond batch processing
- Remote 3D graphics & interactive use, and/or VDI
- Mobile tools for job management

Manage the data

- For IP protection and compliance
- For efficient re-use and retrieval
- In context of PLM, SPDM

Software asset management

- Common tools vs. point solutions
- Shared centralized licensing
- Consolidated business engagement

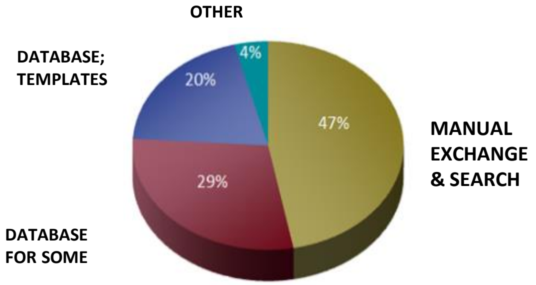

Management of Simulation
Data Source: CPDA

EDR MEDESO

Best-in-Class HPC Performance

- **Parallel improvements** made **release by release**.
- **Outstanding parallel scaling** at an increasing scale of parallelism!
- **Collateral showing** evidence of **HPC performance.**

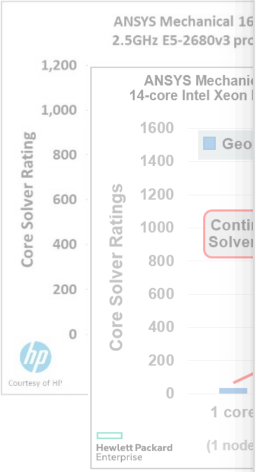

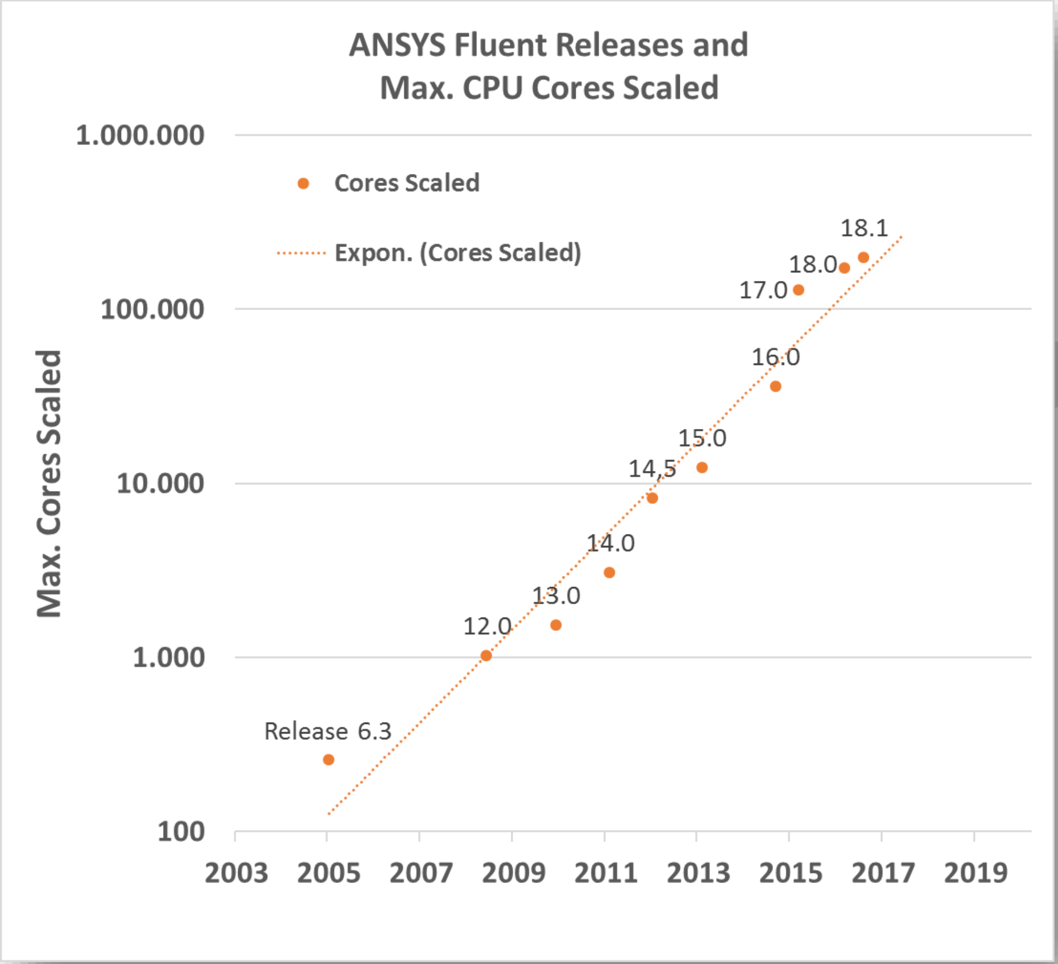

- Be **less constrained** by **hardware limitations** because **'bigger' models** can be **sped up** at your **existing compute capacity!**

EDR MEDESO

Best-in-Class HPC Performance

- **ANSYS Mechanical 2019** demonstrates **65% performance gains** on Intel's latest **Skylake processor** vs. Haswell processor (due to **AVX-512 support**).
- Single-phase simulations in **ANSYS Fluent 19.1** with high AMG solver portion can be **significantly accelerated by GPUs**.

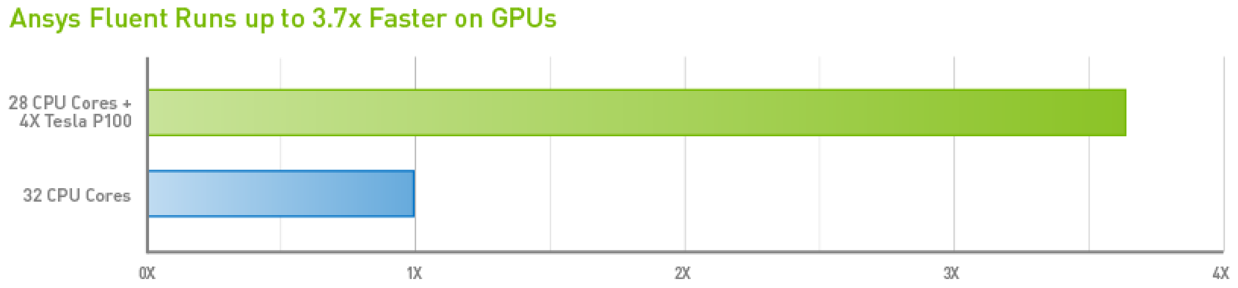

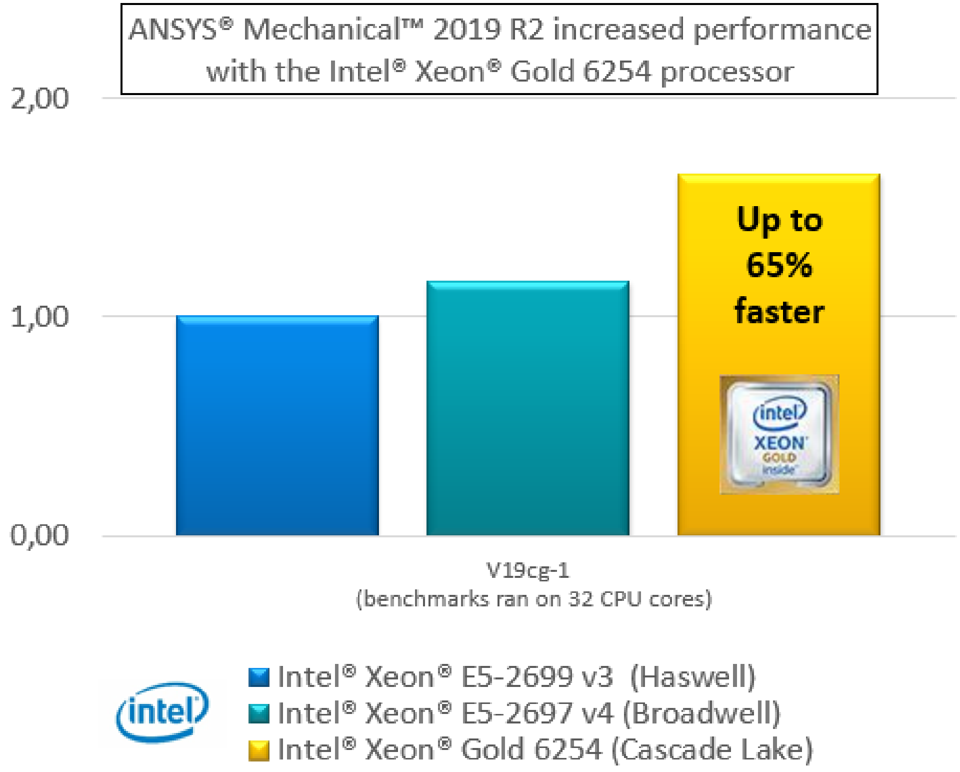

- **Accelerate** your simulation **throughput**.
- Get **more insight** into **product performance.**
- Allow you to **evaluate more** product **design ideas**

Supporting "HPC Resources Anywhere"

ANSYS supports and certifies:

- **ANSYS Cloud** (on-demand HPC in the Azure cloud from within ANSYS Mechanical, Fluent, and Electronics Desktop)
- **Leading remote display software solutions** (VNC, DCV, Exceed onDemand, and Microsoft Remote Desktop)
- **4 proven VDI solutions**
- **Leading job schedulers** (LSF, PBS Professional, UGE/SGE, MOAB/Torque, Microsoft HPC)

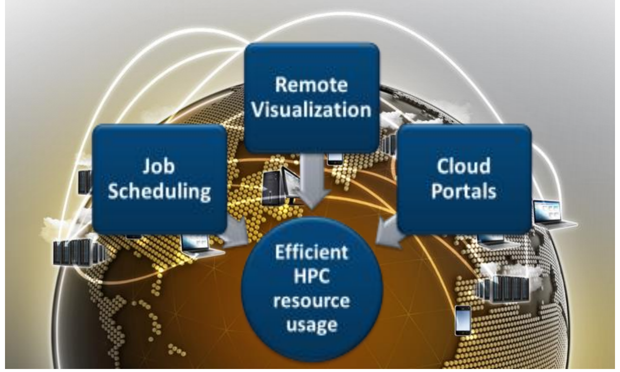

- Easy access to **more powerful HPC resources,** and **simulate models** that **were** simply **impossible** in the past.
- **Collaborate** virtually **from anywhere** with any client device.
- **Increase HPC resource utilization** while **lowering IT support overhead.**
- **Reduce network overload** and **security concerns** by elimination of moving big simulation data sets around!

EDR™ MEDESO

Supporting “HPC Resources Anywhere”

ANSYS Cloud

ANSYS Cloud (launched on Feb 5th, 2019)**:**

- **Easiest** and most **robust way** to access **on-demand HPC on Azure** directly from **within ANSYS Mechanical, Fluent, and Electronics Desktop**
- **Highly optimized** for **ANSYS Mechanical, Fluent,** and **Electronics Desktop** solvers (CFX, LSDyna in H12020)
- **Backed by** the **ANSYS Customer Excellence support team** with decades of **HPC experience** and **application knowledge**

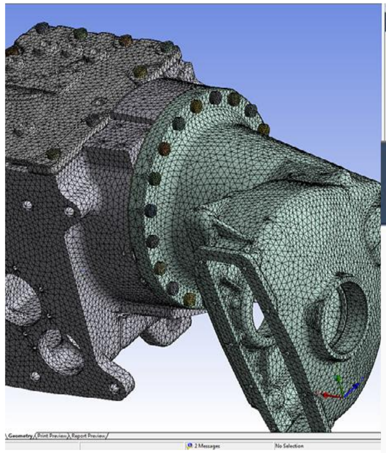

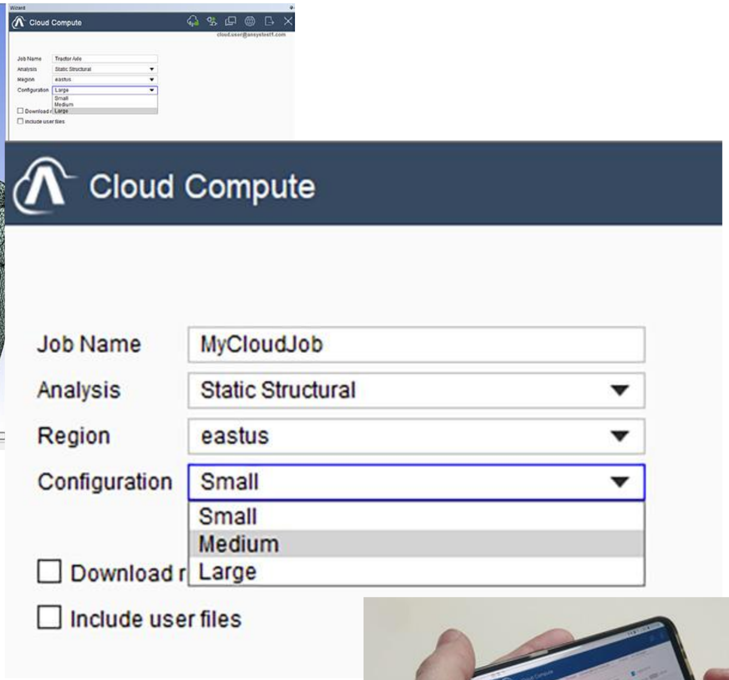

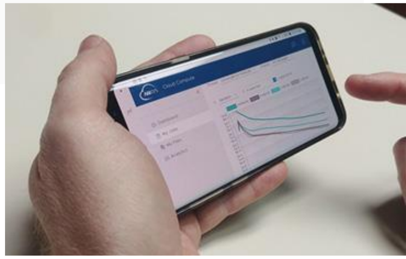

- **Easiest access to HPC** while **avoiding internal IT hurdles** and **delays!**
- **Scale up** and **down** your **computer** and **license capacity** for **variational ANSYS workloads.**

EDR™ MEDESO

Supporting “HPC Resources Anywhere”

Through Cloud-Hosting Partners

Collaborate with CHPs that provide a combination of IT services and infrastructure

- Unique **custom hosting requirements** that **cannot be provided** by **ANSYS Cloud** solution
- **Backed by** a **public cloud** provider **other than Azure**, or **regional datacentre**
- **Leverage a bare-metal,** cloud HPC **infrastructure** that is **suited for** very **computationally demanding** applications

- Through either **existing** or **elastic licenses,** turnkey **access to** either **regional** or **global datacenters!**
- **Scale up** and **down** your **computer** and **license capacity** for **variational workloads.**

EDR™ MEDESO

Remote Display Support

ANSYS 2019 R3 supports the following remote display solutions:

- Nice Desktop Cloud Visualization (DCV) 2017.1
  - Linux server + Linux/Windows client
- OpenText Exceed onDemand 8 SP11
  - Linux server + Linux/Windows client
- OpenText Exceed TurboX 12.0
  - Linux server + Linux/Windows client
- VNC Connect 6.4 (with VirtualGL 2.6)
  - Linux server + Linux/Windows client
- Microsoft Remote Desktop (on Windows cluster)

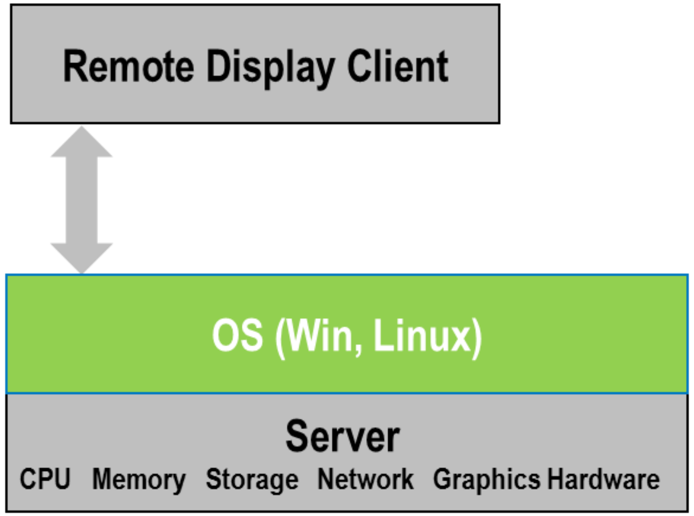

Hardware requirements for remote visualization servers require:

- GPU capable video cards
- large amounts of RAM accessible for multiple user availability when running ANSYS applications and pre/post processing

Supporting “HPC Resources Anywhere”

Virtual Desktop (VDI) Support

Support for virtual GPU

- **for less graphically intensive work** – GPU to be shared between multiple virtual machines (VMs)

GPU pass-through for best performance

- **One GPU per VM**, up to 8 VMs per machine (K1, K2 cards); memory constraints will limit in any case

Supported at 2019 R3:

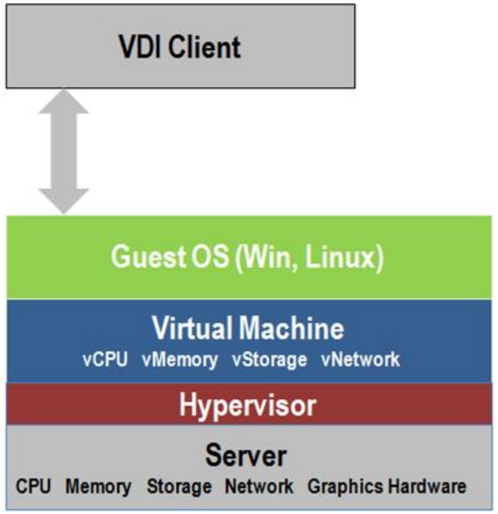

VDI Client	Client OS	Guest OS	Hypervisor	Server Graphics
VMware Horizon view 7.7	Windows	Windows 7, 10, Server 2016, 2019	VMware vSphere ESXi 6.5 U2	NVIDIA GRID (gpu pass-through or vGPU)
Citrix XenDesktop 7.20	Windows	Windows 7, 10, Server 2016	Citrix XenServer 7.6	NVIDIA GRID (gpu pass-through or vGPU)
NICE DCV 2017.4	Windows/Linux	Red Hat 6 and 7, SLES 12, CentOS-7	VMware vSphere ESXi 6.5 U2 / Citrix XenServer 7.1	NVIDIA GRID (gpu pass-through)

EDR MEDESO

Supporting “HPC Resources Anywhere”

Through Scalable Licensing

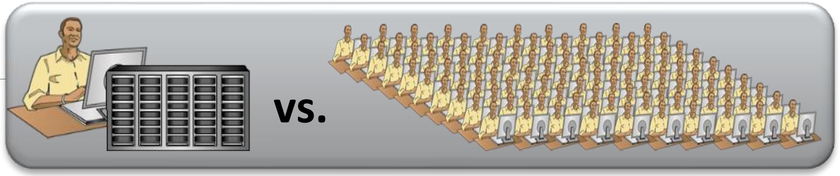

HPC software licensing is related to **customer value** and **computing demands**

- **HPC Pack** offers scalable pricing **for *capability* computing; HPC Workgroup** offers volume pricing **for *capacity* computing**
- **HPC Parametric Pack** and **DSO for *parametric* capacity computing** applied to design exploration
- Usage-based **ANSYS Elastic Units Packs** optimal **for intermittent use** and/or **peak demands of cloud HPC**

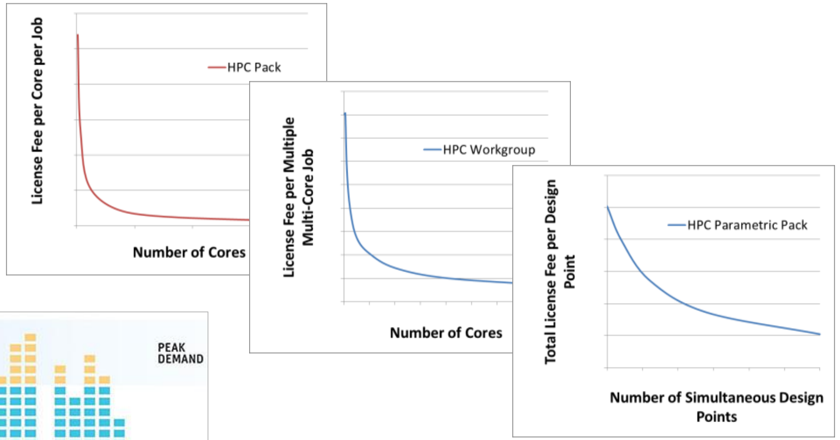

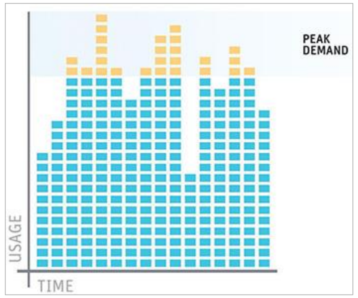

- Running on **2,000 cores instead of 20 cores** incurs a **cost premium of only 1.5X** - and **not** the **100X**!
- Filling up a **1,024- instead of 128-core cluster** with 32-core jobs will **cut the price per job in half!**

EDR™ MEDESO

ANSYS 2019 Benchmark Results and Reference Configurations

EDR MEDESO

HPE Benchmark Environment fOR ANSYS

	HPE Apollo 2000 Gen10 Cluster XL170r Gen10 nodes	HPE Apollo 6000 Gen10 Cluster XL230k Gen10 nodes
Processor Model & Clock Speed	Intel® Xeon® Gold 6142 2.6 GHz	Intel® Xeon® Gold 6242 2.8 GHz
Total Cores per Compute Node	16 cores/Socket (32 cores)	16 cores/Socket (32 cores)
Memory per Node	192GB	192GB
Memory Clock in MHz	2666 MHz	2666 MHz
Network Interconnect	EDR InfiniBand	EDR InfiniBand
Linux OS	Red Hat Enterprise Linux Server release 7.6	Red Hat Enterprise Linux Server release 7.6
Turbo On/Off	Turbo On	Turbo On
Total Cores available	128 nodes/4096 cores	16 nodes/512 cores

ANSYS Mechanical 2019 R1 using Intel Xeon Gold 6142 2.6 GHz

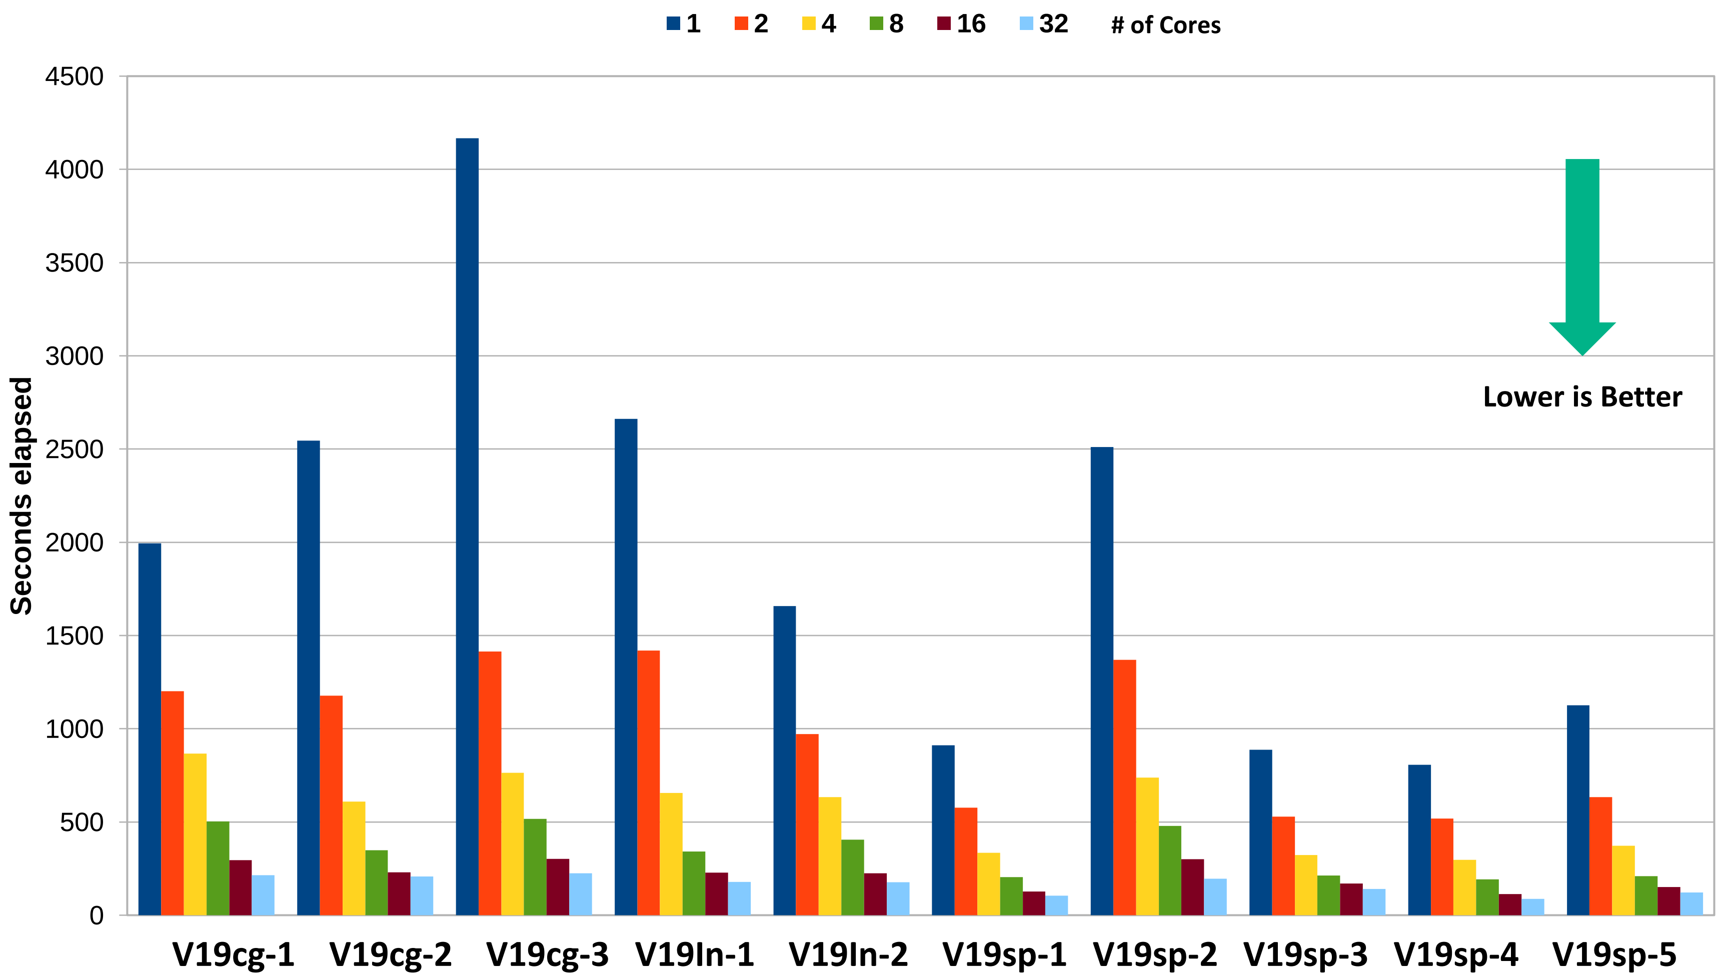

- Power Supply Module (V19cg-1)
- Tractor Rear Axle (V19cg-2)
- Engine Block (V19cg-3)
- Gear Box (V19ln-1)
- Radial Impeller (V19ln-2)
- Peltier Cooling Block (V19sp-1)
- Semi-Submersible (V19sp-2)
- Speaker (V19sp-3)
- Turbine (V19sp-4)
- BGA (V19sp-5)

EDR™ MEDESO

ANSYS Mechanical 2019 R3 - up to 4 nodes (128 cores) Intel Xeon Gold 6242 2.6 GHz

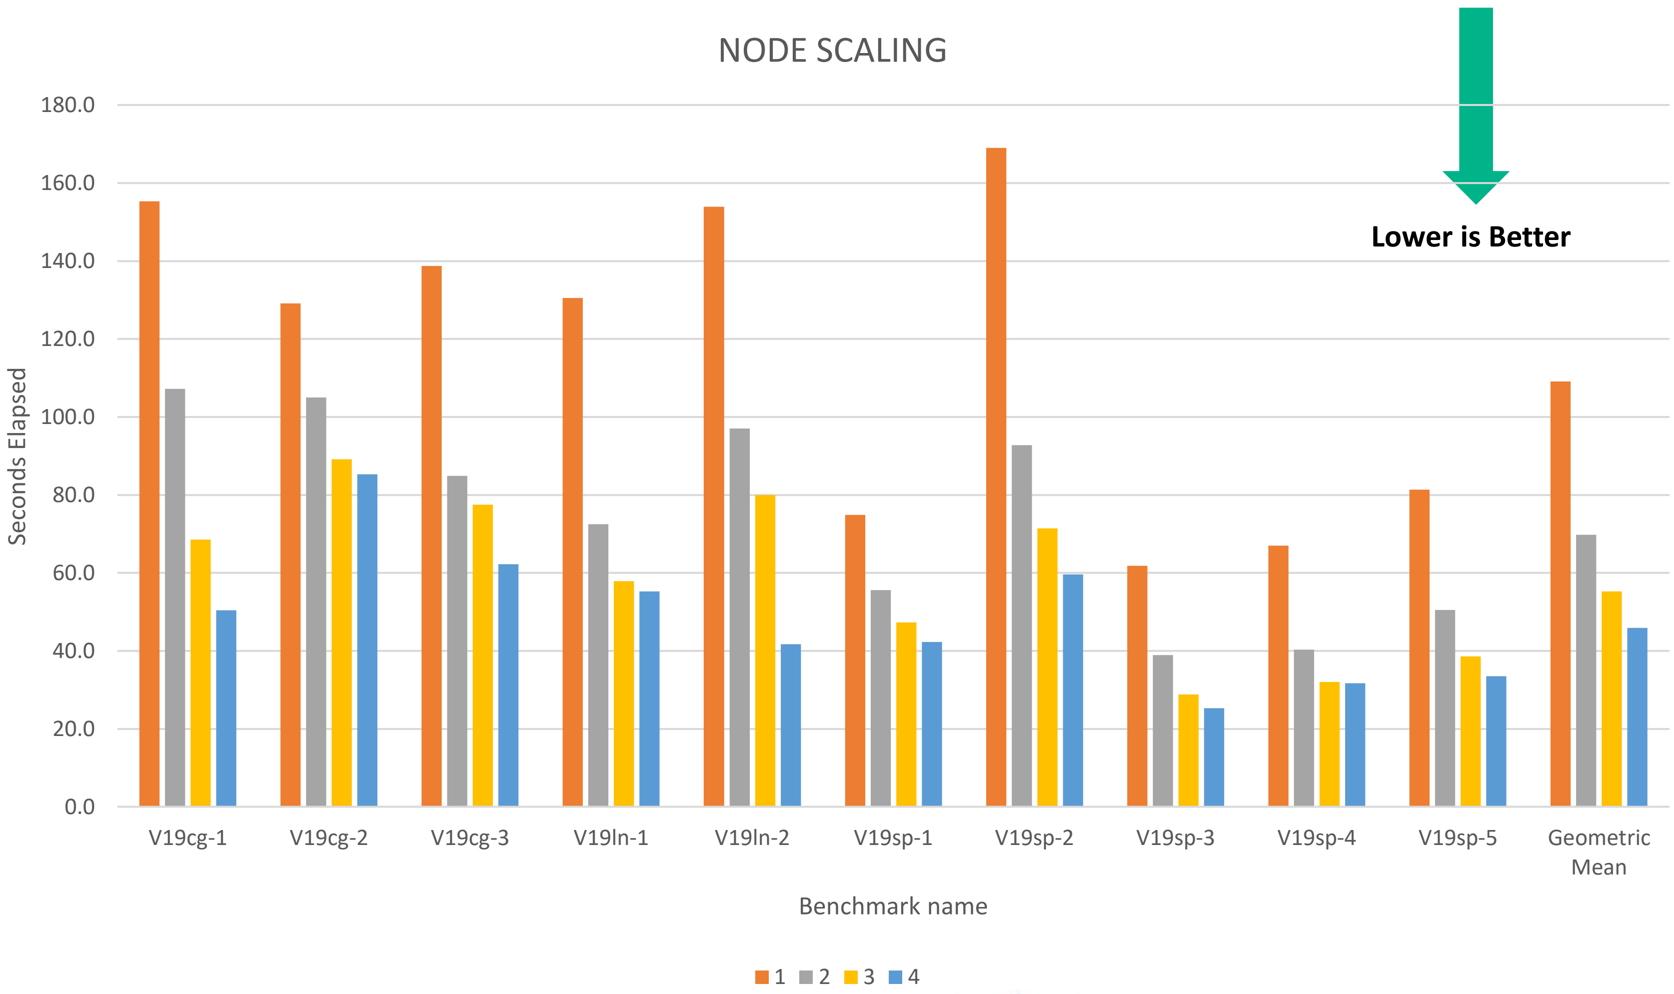

- Power Supply Module (V19cg-1)
- Tractor Rear Axle (V19cg-2)
- Engine Block (V19cg-3)
- Gear Box (V19ln-1)
- Radial Impeller (V19ln-2)
- Peltier Cooling Block (V19sp-1)
- Semi-Submersible (V19sp-2)
- Speaker (V19sp-3)
- Turbine (V19sp-4)
- BGA (V19sp-5)

EDR MEDESO

ANSYS Mechanical 2019 up to 4 nodes –
2019R1 Xeon Gold 6142 (SKL) - 2019R3 XEON GOLD 6242 (CCL)

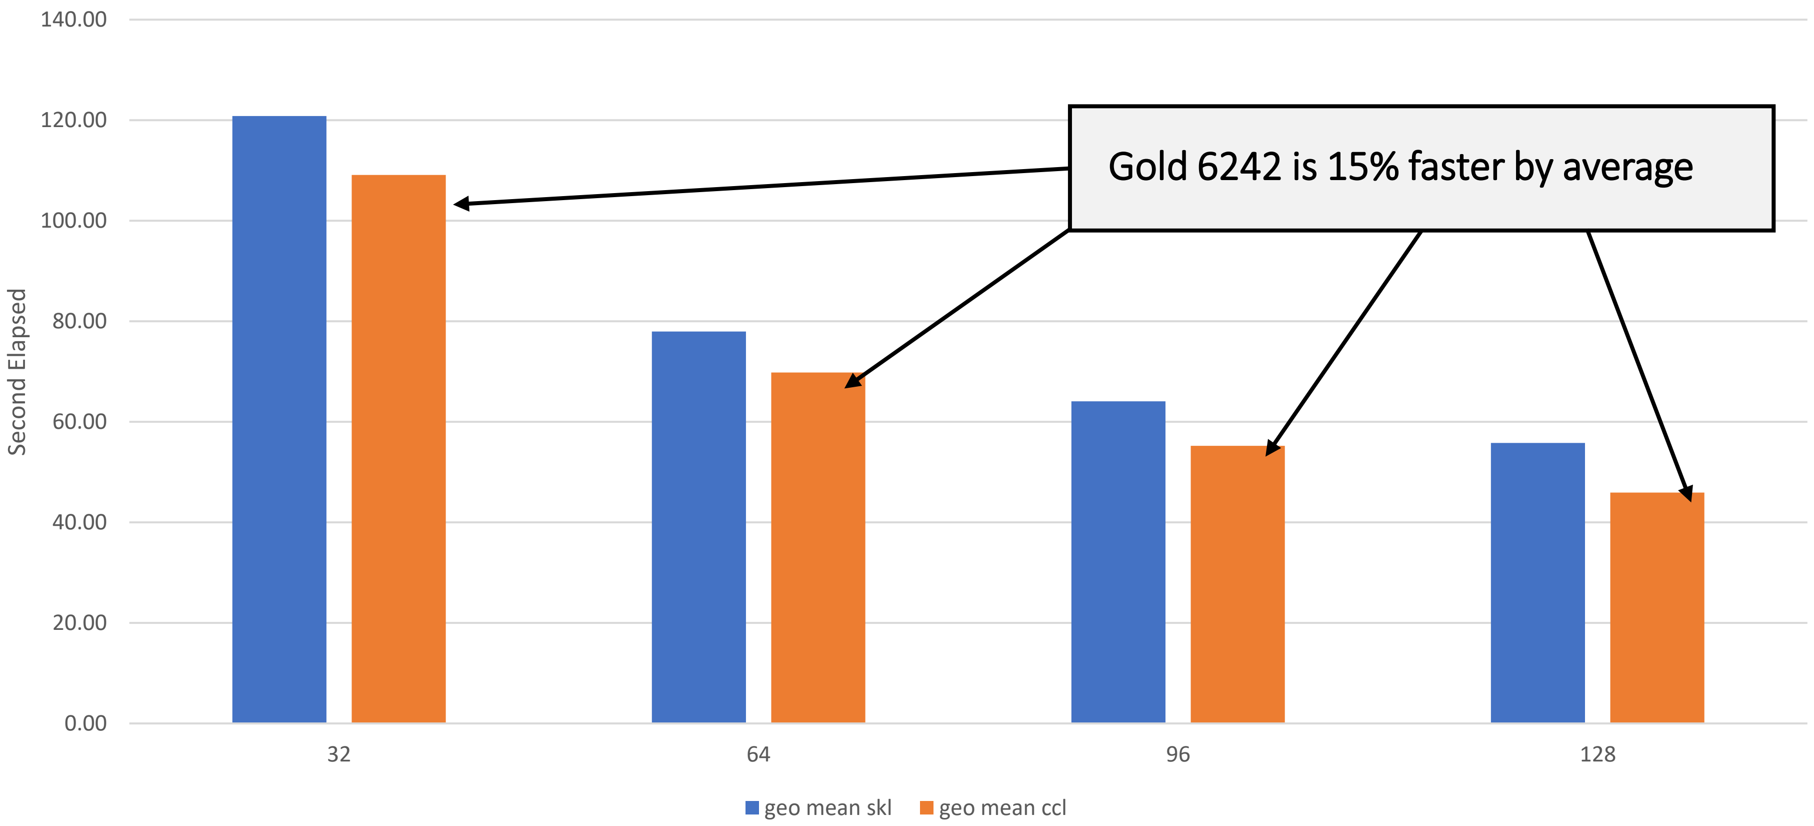

Lower is Better

ANSYS Fluent 2019 R1 Speedup - Single Intel Xeon Gold 6142 2.6 GHz processor

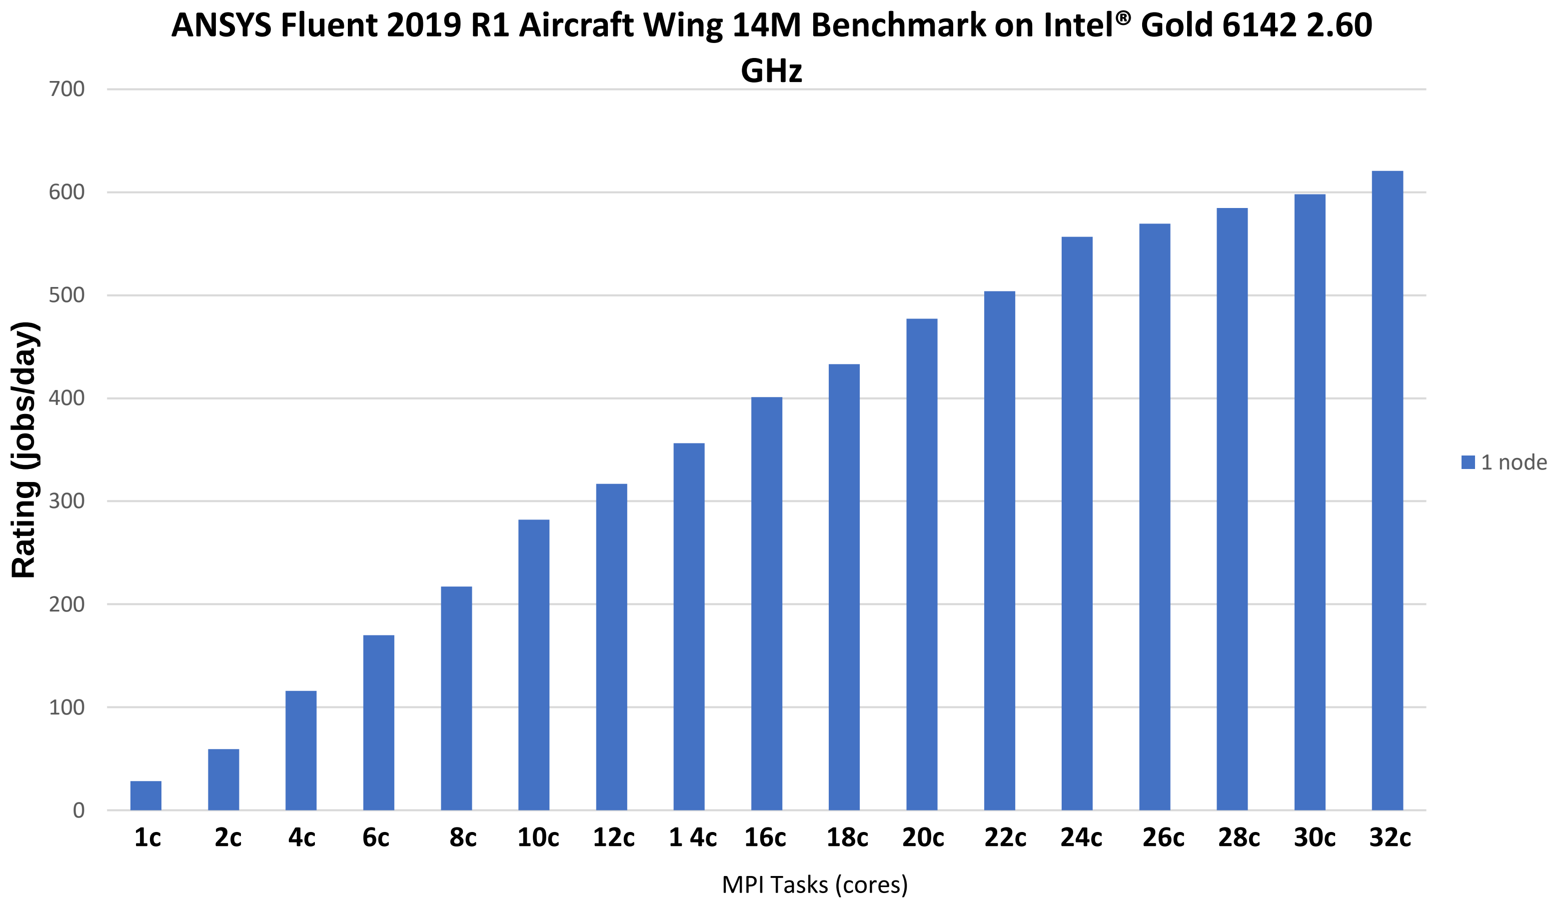

Higher is Better

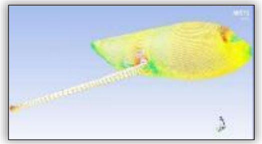

ANSYS Fluent Standard Benchmark Aircraft Wing 14M Cells

EDR MEDESO

ANSYS Fluent 2019 Performance comparison:
Intel Xeon Gold 6142 2.6 GHz vs. Intel Xeon Gold 6242 2.8 GHz

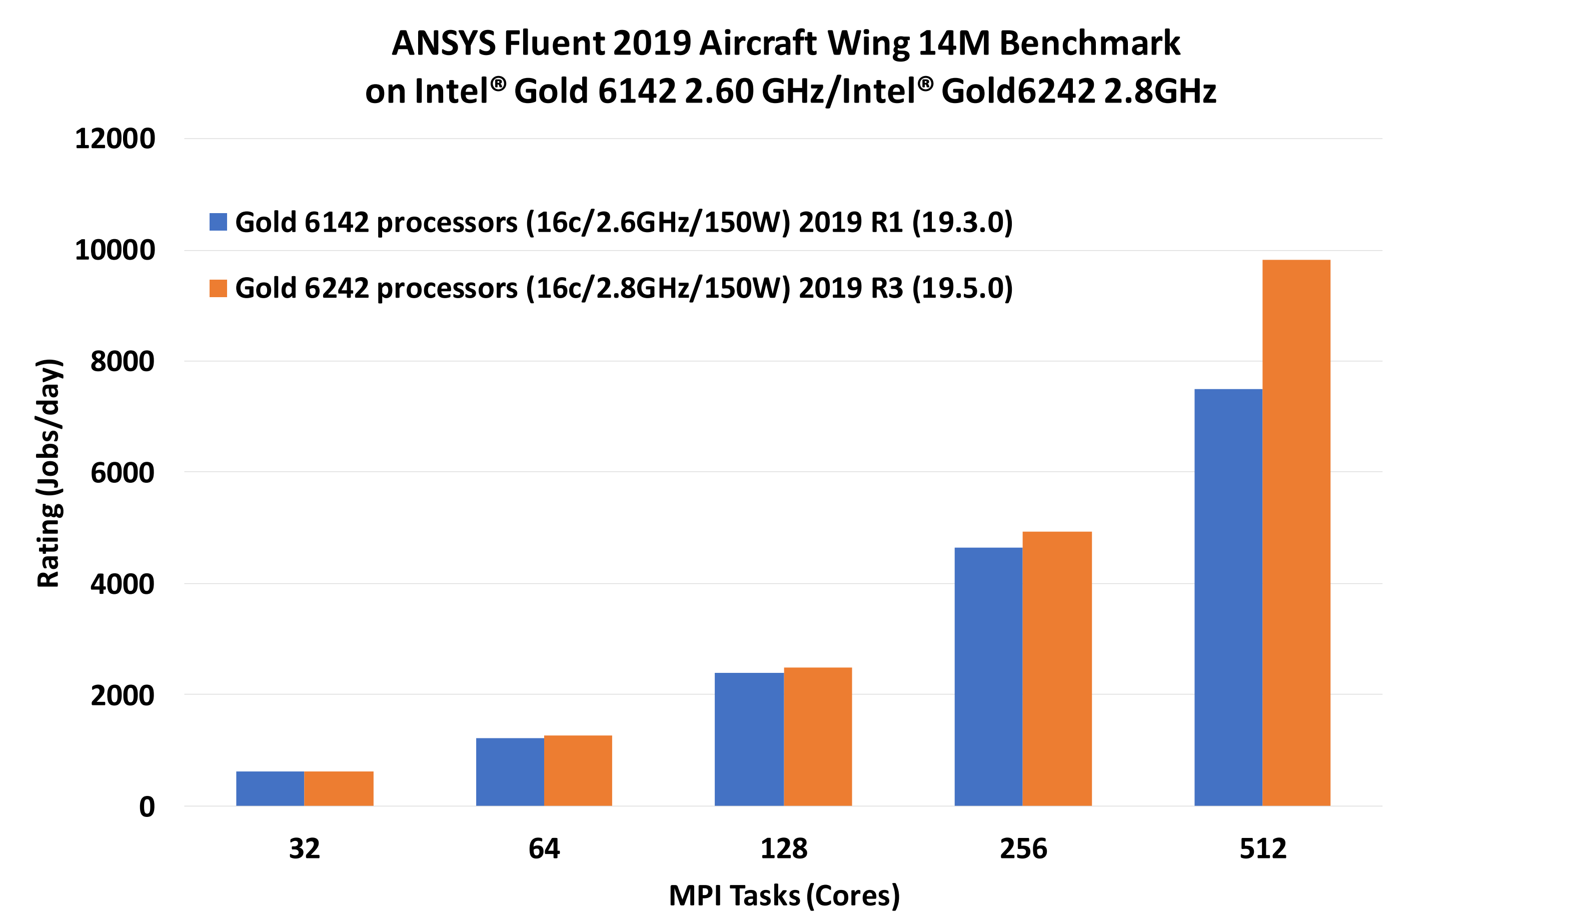

Higher is Better

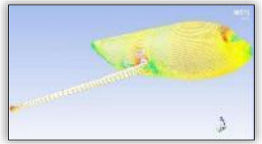

ANSYS Fluent Standard Benchmark Aircraft Wing 14M Cells

EDR™ MEDESO

ANSYS Fluent 2019 Performance comparison:
Intel Xeon Gold 6142 2.6 GHz vs. Intel Xeon Gold 6242 2.8 GHz

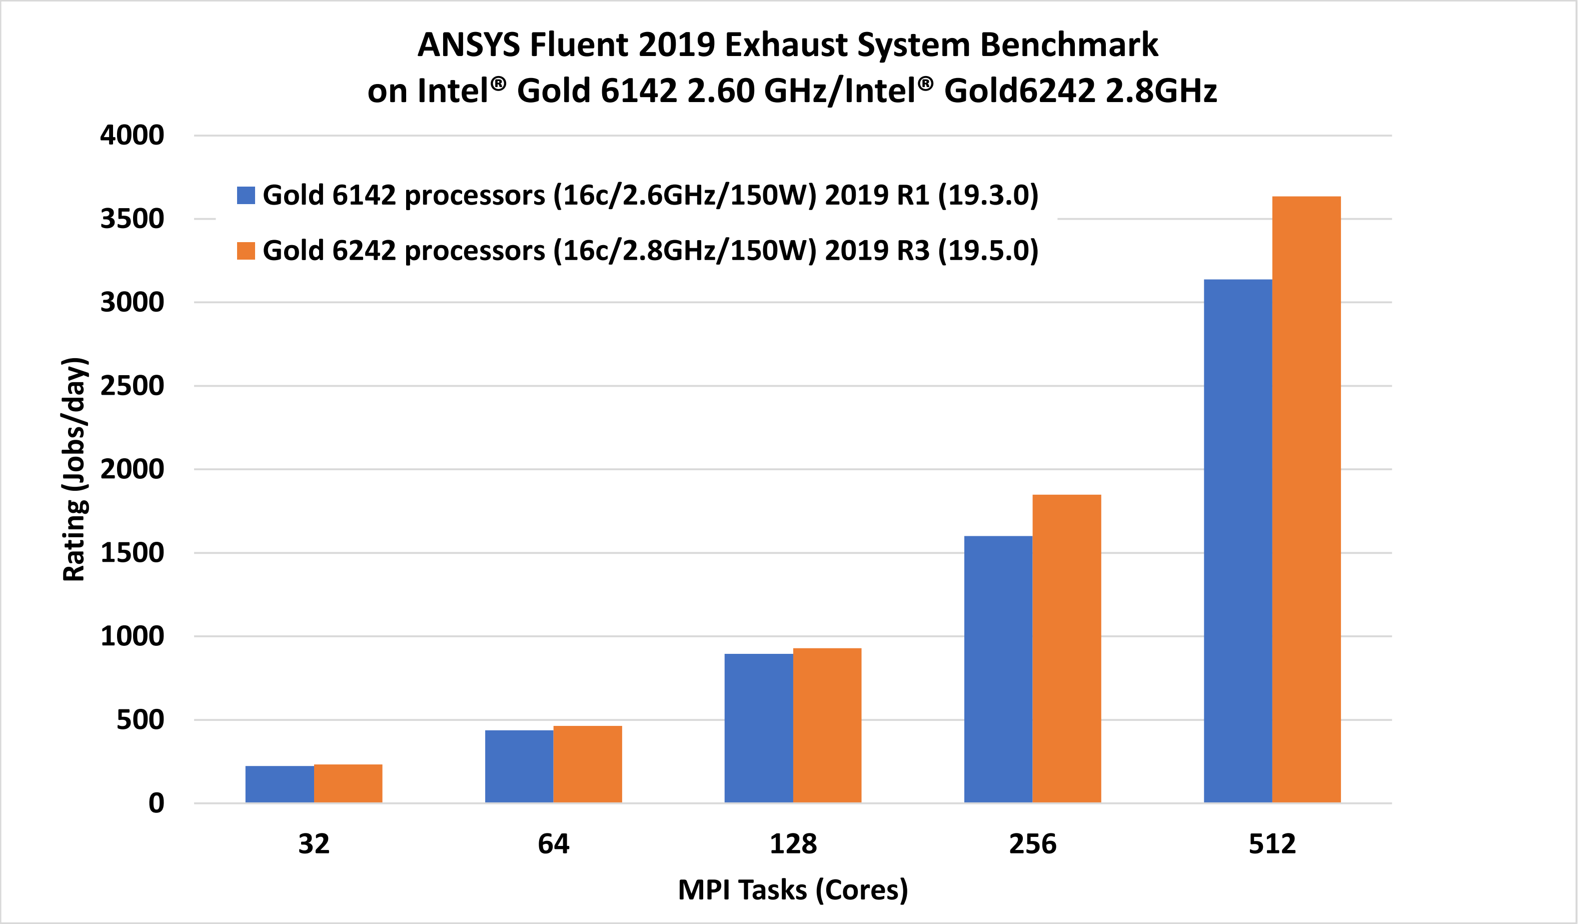

ANSYS Fluent Standard Benchmark Exhaust System 33M Cells

EDR™ MEDESO

ANSYS Fluent 2019 Performance comparison:
Intel Xeon Gold 6142 2.6 GHz vs. Intel Xeon Gold 6242 2.8 GHz

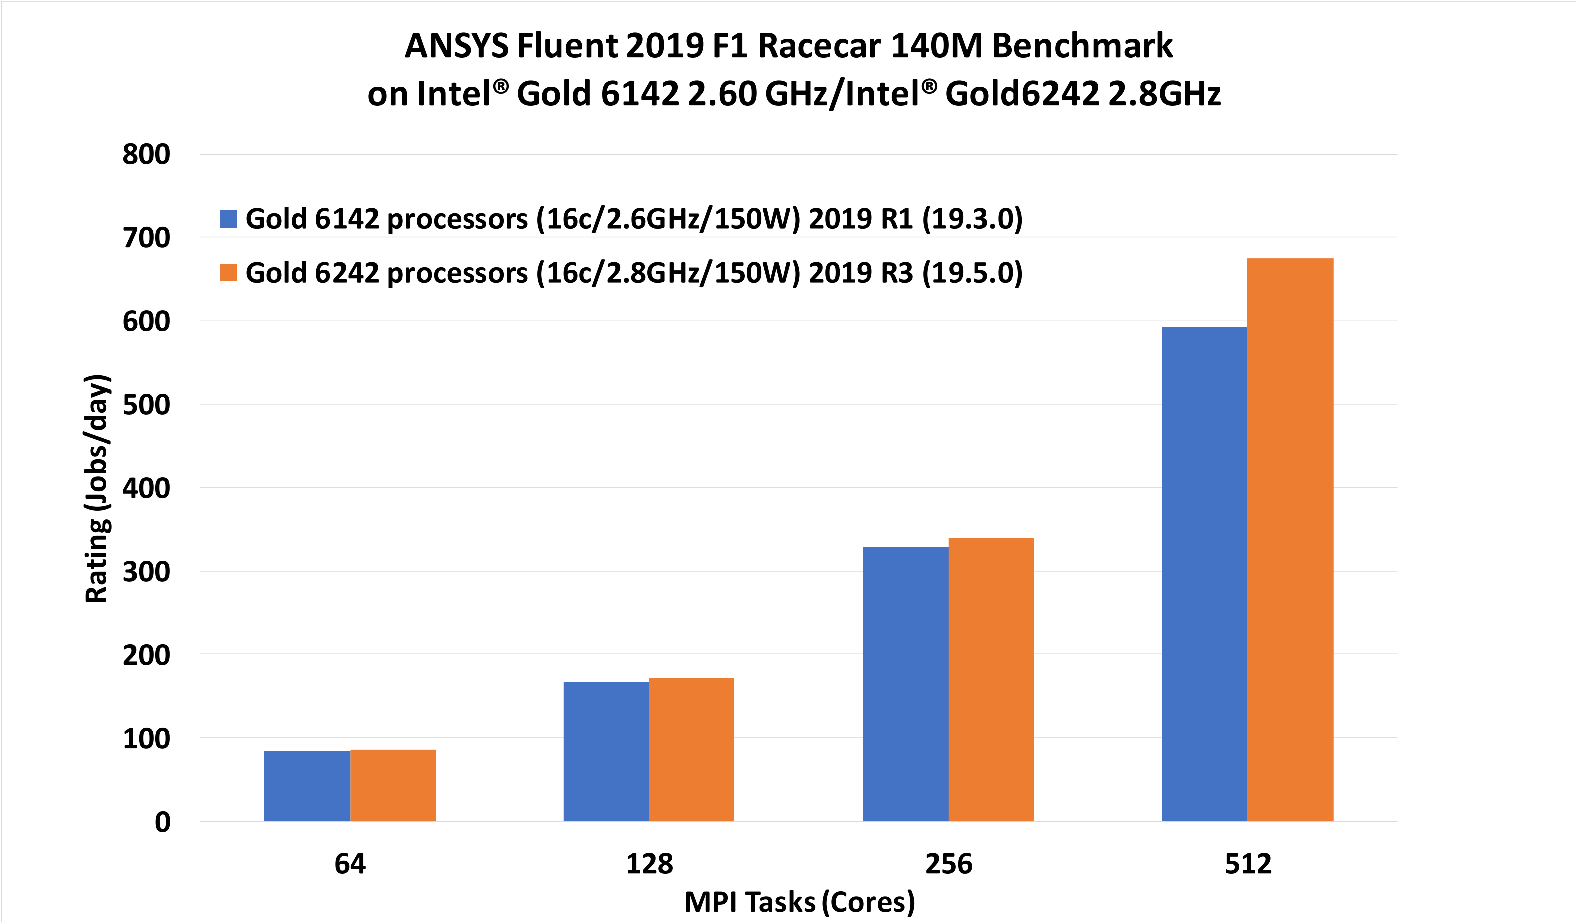

Higher is Better

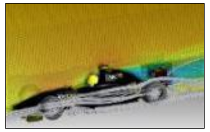

ANSYS Fluent Standard Benchmark F1 Racecar 140M Cells

ANSYS CFX 2019R3 scaling running the LeMans Car Model

	HPE SGI 8600 Cluster	HPE Apollo 6000 Gen10 Cluster XL170r Gen10 nodes	SGI ICE XA (legacy system)
Processor Model & Clock Speed	Intel® Xeon® Gold 6154 3.0 GHz	Intel® Xeon® Platinum 8268 2.90 GHz	Intel® Xeon® E5-2690v4 2.6 GHz
Total Cores per Compute Node	18 cores/Socket (36 cores)	16 cores/Socket (32 cores)	14 cores/Socket (28 cores)
Memory per Node	192GB	192GB	128GB
Memory Clock in MHz	2666 MHz DDR4	2933 MHz DDR4	2400 MHz DDR4
Network Interconnect	Intel Omni-Path OPA	Intel Omni-Path OPA	EDR InfiniBand
Linux OS	Red Hat Enterprise Linux Server release 7.6	Red Hat Enterprise Linux Server release 7.6	SLES 11 SP3
Turbo On/Off	Turbo On	Turbo On	Turbo On
Total Cores available	288 nodes/10368 cores	288 nodes/13824 cores	128 nodes/3,584

ANSYS CFX 2019 R3 scaling up to 8 nodes of various Intel Xeon processors

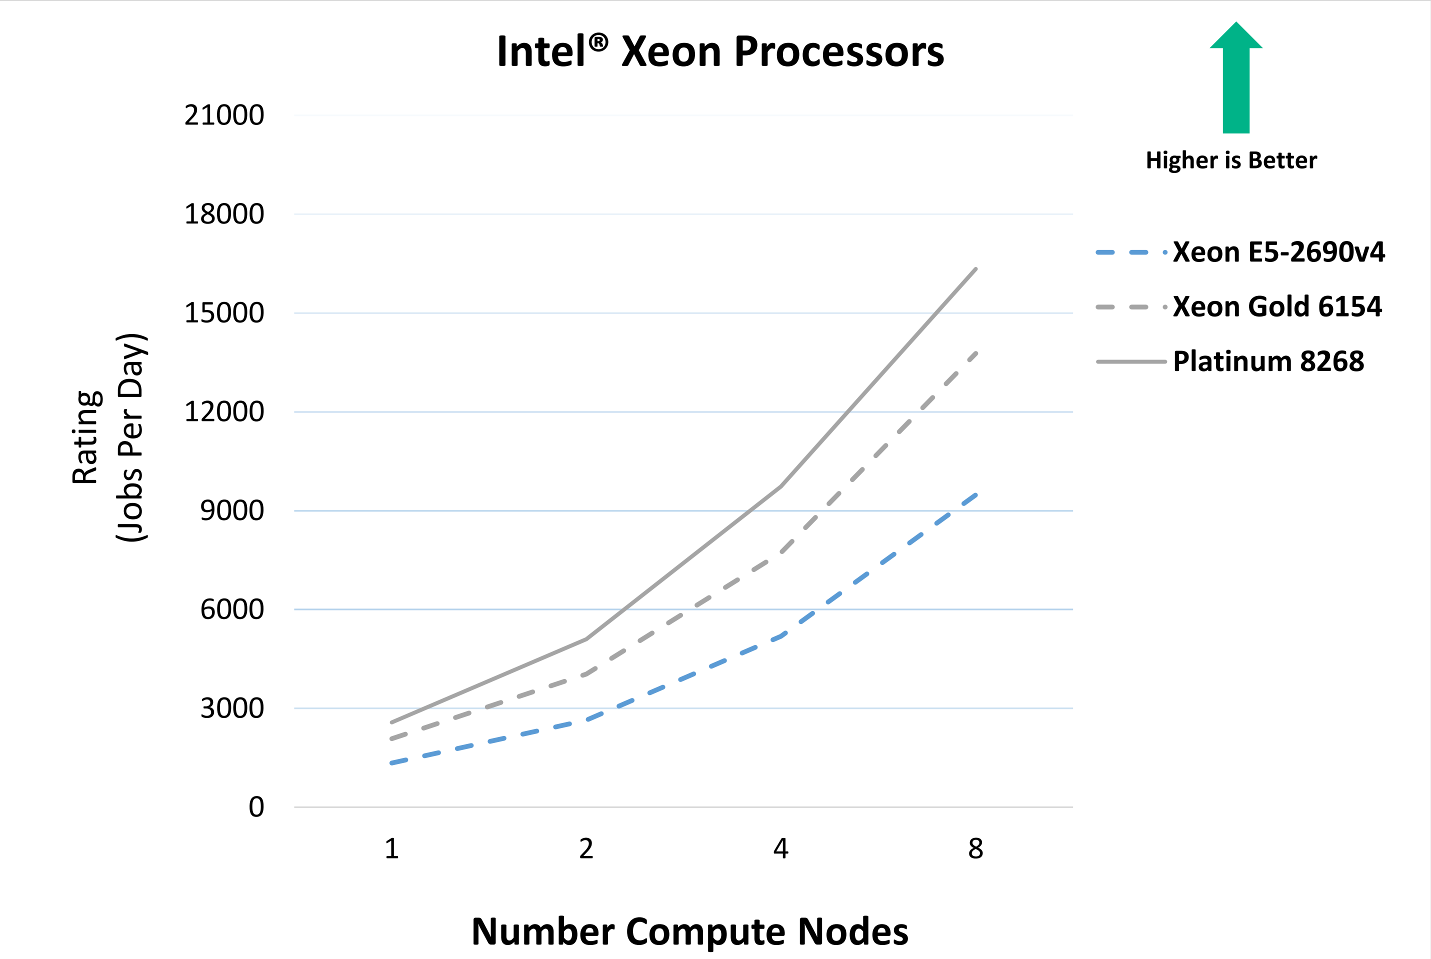

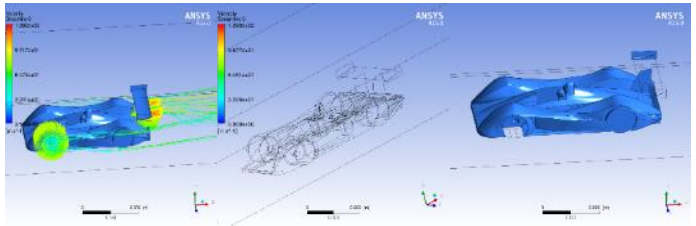

LeMans Car Model

External flow over a LeMans Car. The case has approximately 1.8 million nodes (10 million elements, all tetrahedral), and solves compressible fluid flow with heat transfer using the k-epsilon turbulence model.

EDR™ MEDESO

High Frequency Simulations with HPC

- HF simulations benefit from high performance computing
  - Thump rule: More CPU cores & More memory → Faster
  - Benefits of the parallel computing depend heavily on the type of calculation methods and the way the simulation model boundaries and variables have been set.
- HF simulations have two major steps with different performance
  - Initial mesh and automatic adaptive mesh creation
  - Post processing steps – frequency sweeps, optimisation, tuning, Design of Experiments, …

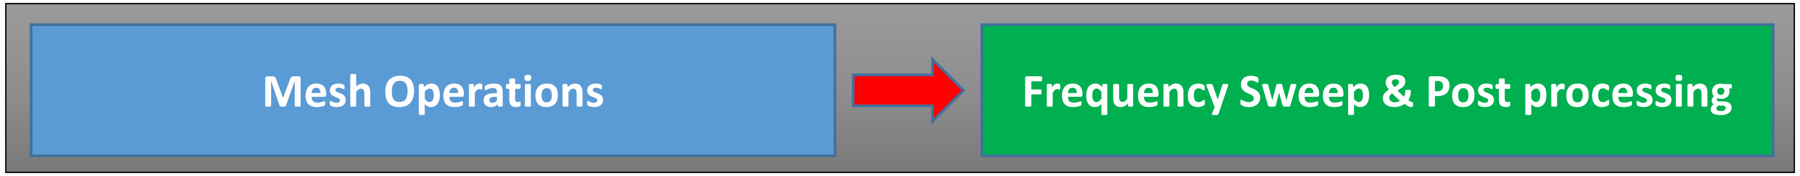

High Frequency (HF) Simulations with HPC

Mesh creation - HFSS

- Initial mesh steps use mainly <u>one CPU core</u> <-> good to have a very fast single CPU core performance and a lot of RAM memory
- Meshing phase uses multiple cores with certain steps
  - → speed improvement with multicore CPU (4 … ~40 cores)
- Mesh creation require RAM memory
  - → the more cells the more memory… 2 M cells → 128 GB Ok
- Iterative mesh creation can take 20%...90% of the total simulation time

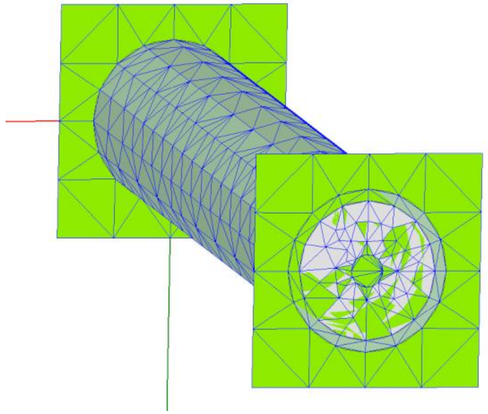

Task	Real Time	CPU Time	Memory	
Start				Time: 03/10/2020 11:31:22; Host: 1990LTW075; Processo
				Executing from C:\Program Files\AnsysEM\AnsysEM20.1
				HPC Enabled
Mesh (surface, seed)	00:00:04	00:00:04	185 M	68602 triangles
				Length1
Adaptive Pass #1				
Machine Configuration				1990LTW075 Using 8 core(s); 33406240 K aval. mem.
Solver setup	00:00:21	00:02:25	606 M	8 core(s)
Matrix solution	00:00:03	00:00:21	818 M	8 core(s)

EDR MEDESO

High Frequency (HF) Simulations with HPC

- The initial mesh and iterative mesh step must be done for each;
  a. physical 3D CAD model variation → as many mesh steps as physical setups with varying shapes
  b. defined mesh solution frequency in a multifrequency simulation
  c. 3D setup where material changes from dielectric to conductor
- Individual mesh phases can run in parallel
  - Linear speed up with HPC
- Note that each individual mesh calculation requires certain amount of RAM. The total amount of memory needed depends on the number of parallel operations.
  - Possible to distribute calculations to several workstations for extra resources.

High Frequency (HF) Simulations with HPC

Frequency Sweep

- HPC enables more CPU cores for frequency point calculations
- HPC speeds up calculations nearly linearly as frequency points are calculated independently of each other. However, each frequency point calculation requires a certain amount of RAM → the amount of total memory can limit scaling; >2 M cells require >256 GB
- It is possible to distribute frequency point calculations to multiple PCs.
- The higher mesh count, the more cores are needed to calculate frequency sweep points in parallel.
  - 34 cores & 1.7 M cells → 7 freq. points

Time: 03/20/2020 08:03:50; Host: TAM03; Processor: 40; OS: NT 10.0; HFSS 2020.1.0
Executing from C:\Program Files\AnsysEM\AnsysEM20.1\Win64\HFSSCOMENGINE.exe
HPC Enabled
Allow off core: True
Using automatic settings
Machines:
TAM03.connect365.net [191.2 GB]: RAM Limit: 85.000000%, 34 cores
Solution Basis Order: 1
Discrete HFSS Frequency Sweep, Solving Distributed - up to 7 frequencies in parallel
From 1MHz to 1.5GHz, 251 Frequencies
HPC Enabled

Maxwell with TDM (Time decomposition method)

- Maxwell 2019 R3
- Maxwell design: 3D Electric Motor
- Number of elements(Mesh) = 59 081 elements
- Time step = 200 steps

Scale up in Maxwell with TDM

Simulation	Time	Speed up
1 task 1 core	08.27.45	-
1 task 4 cores	02.42.40	3.12
1 task 8 cores	02.07.19	4
4 tasks 36 cores	01.49.41	4.62
8 tasks 36 cores	00.59.59	8.47
18 tasks 36 cores	00.42.37	11.81
36 tasks 36 cores	00.31.44	15.88
8 tasks 72 cores	00.58.08	8.76
108 tasks 108 cores	00.17.34	28.22
216 tasks 216 cores	00.14.47	33.87

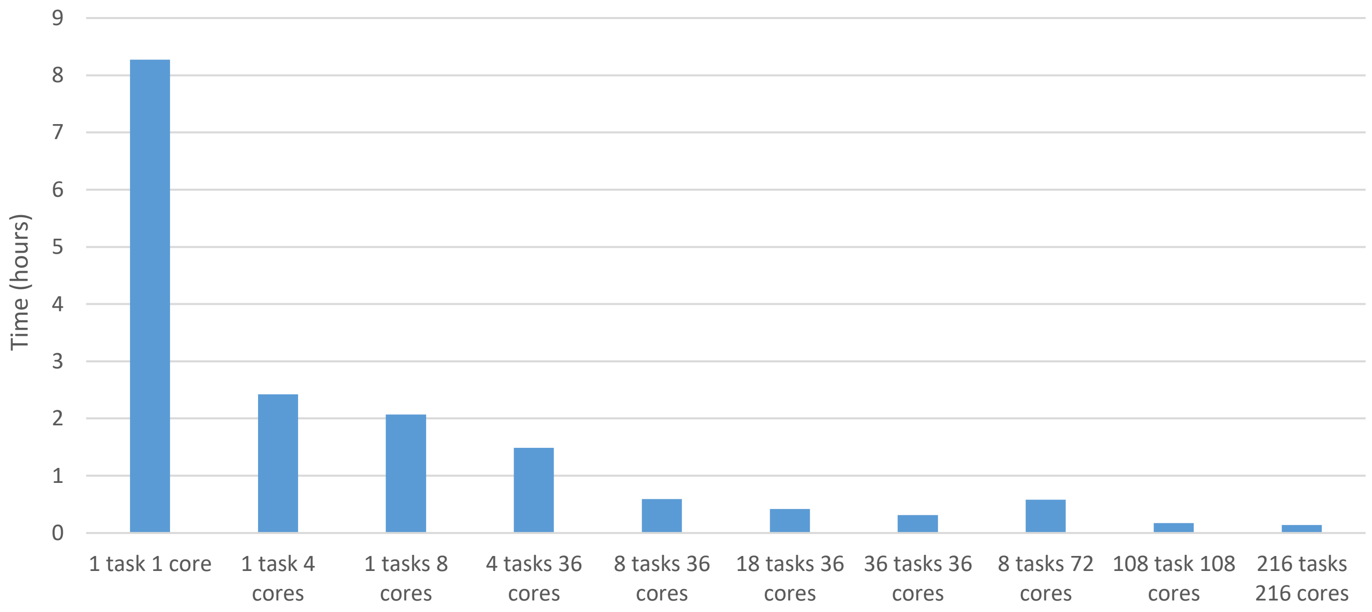

Can se that we get a speed up by increasing the number of cores. This speed up increases even more when the number of tasks (time steps) running simultaneously is increased.

It can be observed that it is more efficient to have several tasks running in parallel, that to allocate all cores to one task.

EDR MEDESO

Questions?

Klas.Johansson@edrmedeso.com