Extreme scalability in CAE ISV Applications

Greg Clifford
Manufacturing Segment Manager
clifford@cray.com

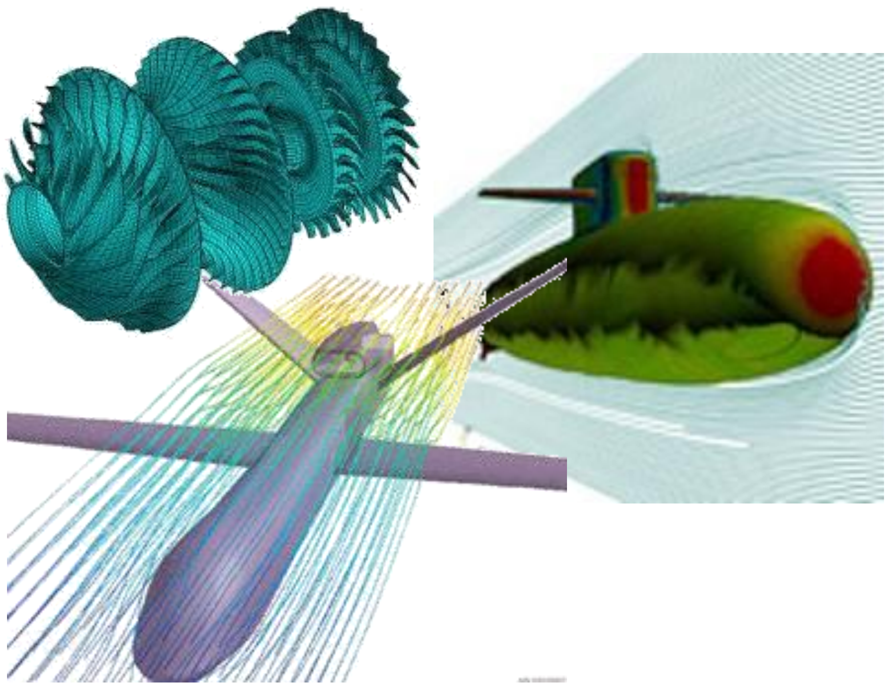

Introduction

- ISV codes dominate the CAE commercial workload
- Many large manufacturing companies have >>10,000 cores HPC systems
- Even for large organizations very few jobs use more than 128 MPI ranks
- There is a huge discrepancy between the scalability in production at large HPC centers and the commercial CAE environment

Can ISV applications efficiently scale to 1000's of MPI ranks?

Spoiler alert: the answer is YES!

Often the full power available is not being leveraged

Is there a business case for scalable CAE applications?

Seismic processing Compute requirements

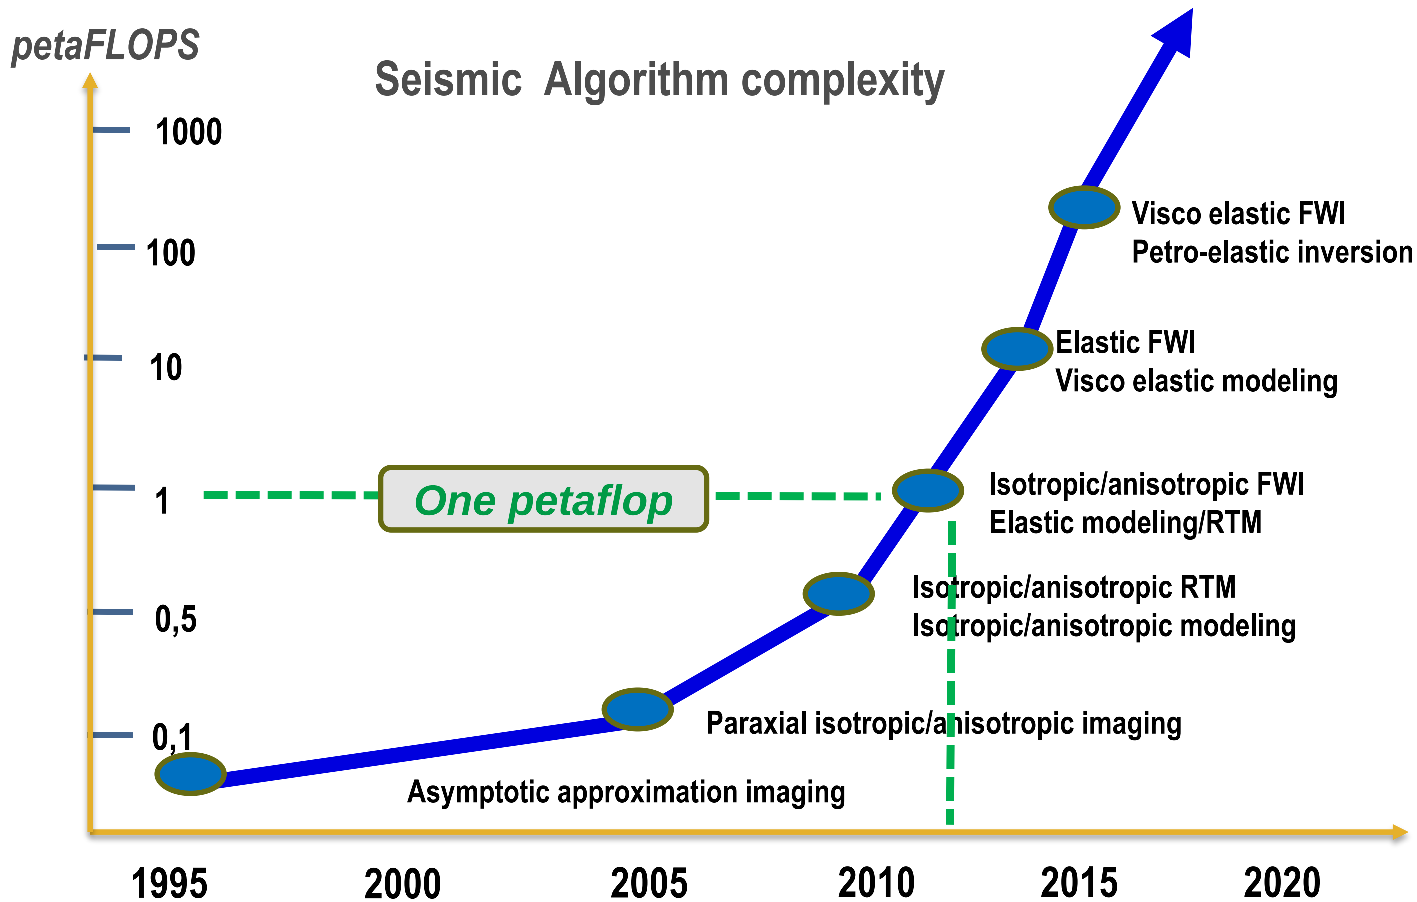

A petaflop scale system is required to deliver the capability to move to a new level of seismic imaging.

Breaking News: April 18, 2012

ExxonMobil

ROSNEFT

ExxonMobil and Rosneft…could invest over ***$500 billion*** in a joint venture to explore for and produce oil in the Arctic and the Black Sea…

…recoverable hydrocarbon reserves at the three key Arctic fields are estimated at ***85 billion barrels***

by the Associate Press

Petascale Seismic Processing: A Business Case

Compute & data requirements for seismic processing are huge

- Wide demands on processing from data acquisition to seismic to res sim
- Petaflop scale systems required for state-of the art processing
- Petabytes of capacity and terabytes of bandwidth from I/O

s

- A sing … er
- "Resta …

The Oil & Gas industry has typically led the way on new HPC hardware technology in the commercial sector

the demand for "getting it right" goes up

- This is the class of simulation that drives real petascale capability computing
- You can do capacity on capability systems but not vice versa – risk mitigation

CAE trends driving HPC requirements

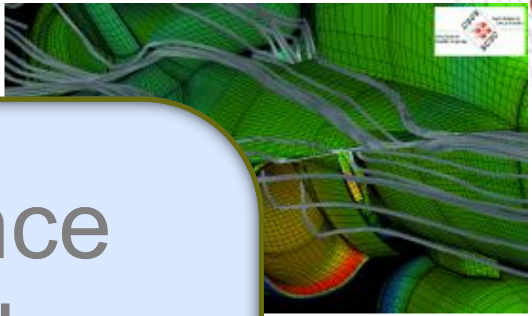

- **"Extreme Fidelity"***
  - Enhanced
    e.g. 1 Bil
  - Large mo
    cores
- **Design op**
  - Many sim
    design sp
  - Multiple r
- **Robust de**
  - Looking for the "best solution"

"Future performance depends on highly scalable parallel software"

ANSYS®

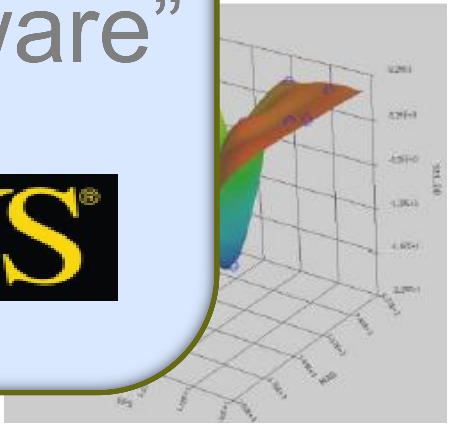

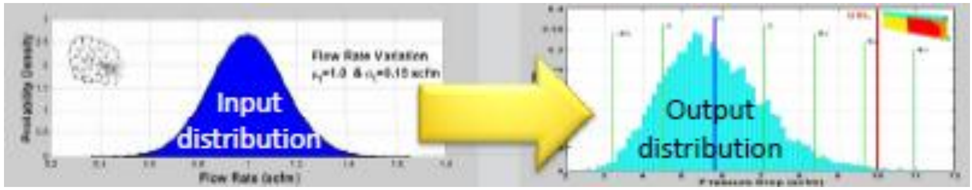

* ref: ANSYS CFD presentation

Compute requirements in CAE

"Simulation allows engineers to know, not guess – but only if IT can deliver dramatically scaled up infrastructure for mega simulations….
1000's of cores per mega simulation" ANSYS

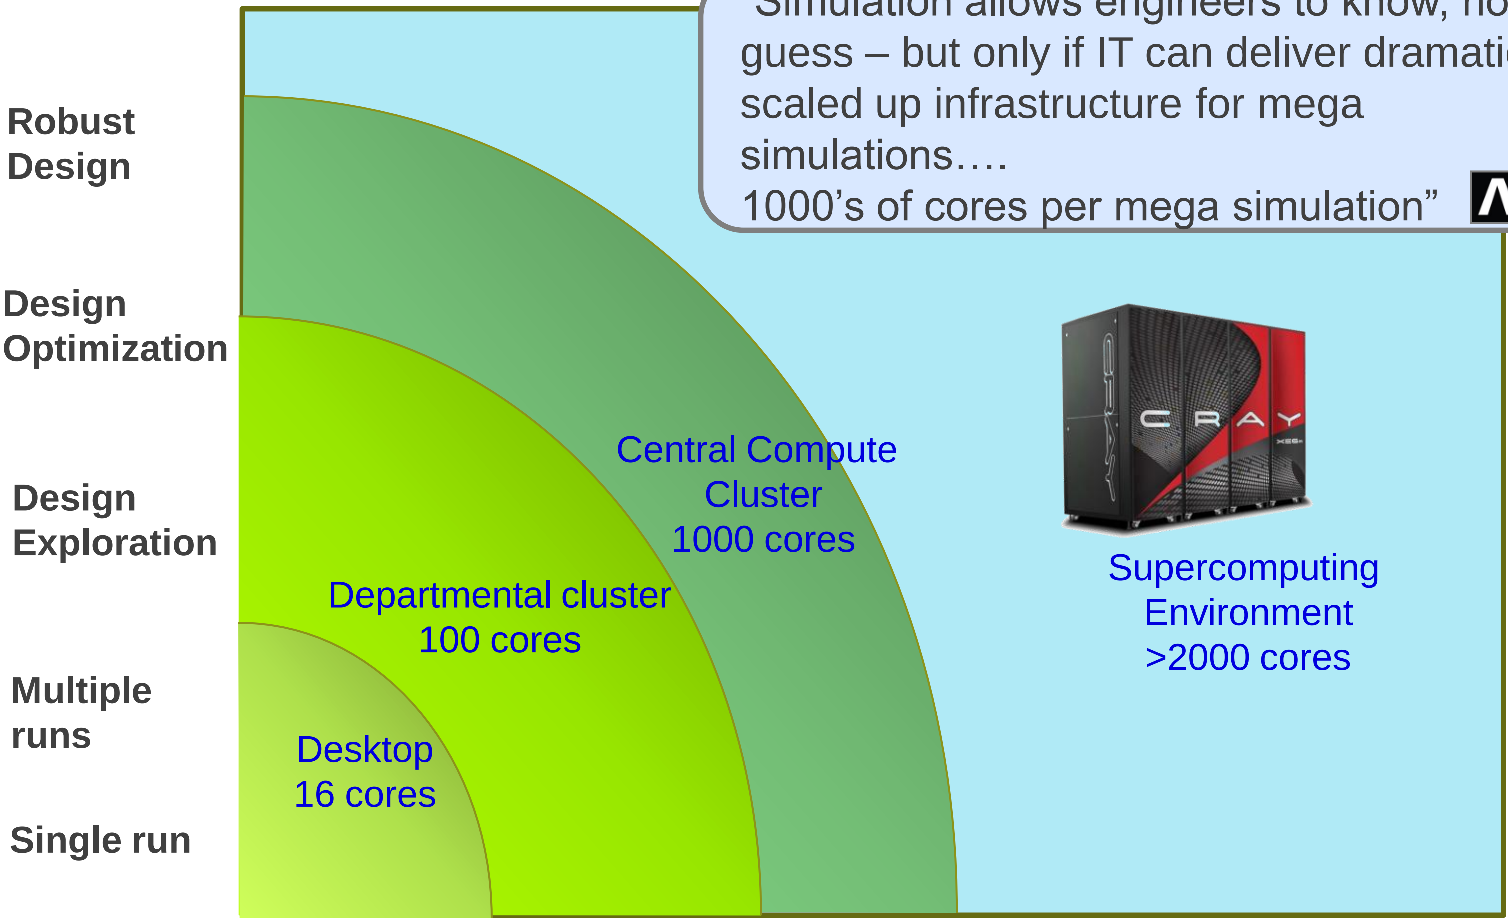

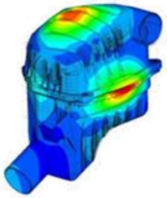

Simulation Fidelity

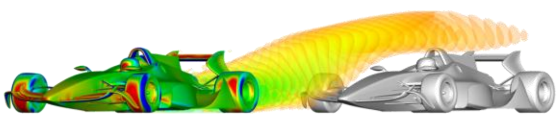

CAE Application Workload

Basically the same ISV codes used across all industries

- Impact/Crash Apps
  - ABAQUS explicit
  - LS-DYNA
  - PAM-CRASH
  - RADIOSS
- CFD Apps
  - CFD++
  - ANSYS Fluent
  - PowerFLOW
  - STAR-CCM+
  - “in-house”
- Structures Apps
  - ABAQUS implicit
  - ANSYS Mechanical
  - MSC.Nastran
- Electromagnetic Apps
  - “in-house” (classified)
  - ANSYS HFSS

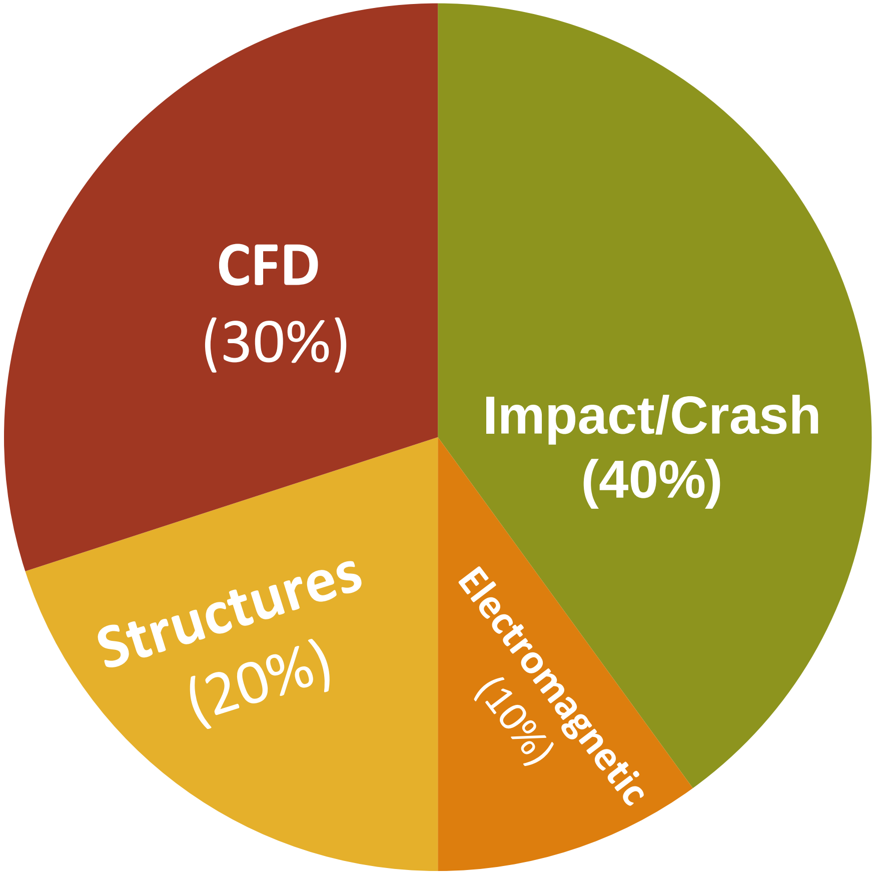

Vast majority of large simulations are MPI parallel

Is the extreme scaling technology ready for production CAE environments?

Brief history of HPC technology in high end environments

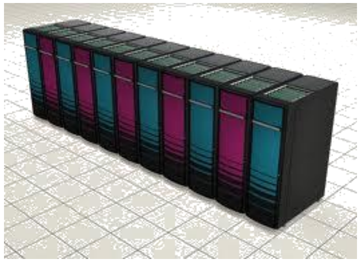

c. 2007, extreme scalability
Proprietary interconnect
1000's cores
Requires "end-to-end parallel"

c. 1998, low density, slow interconnect
"Linux cluster", MPI Parallel
100's of "cores"
Major code restructuring

c. 1983
Cray X-MP, SMP parallel
8 Processors
Compiler directives for key kernels

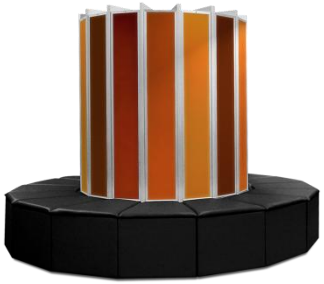

c.1978
Cray-1, Vector processing
1 Processor
Automated vectorization in the compiler

Propagation of HPC to commercial CAE

Early adoption

Common in Industry

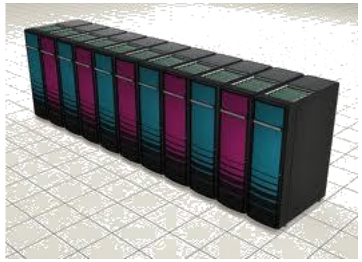

c. 2007, Extreme scalability
Proprietary interconnect
1000's cores
Requires "end-to-end parallel"

c. 2013
Cray XE6
driving apps:
CFD, CEM, ???

c. 1998,
MPI Parallel
"Linux cluster",
low density, slow interconnect
~100 MPI ranks

c. 2003,
high density,
fast interconnect
Crash & CFD

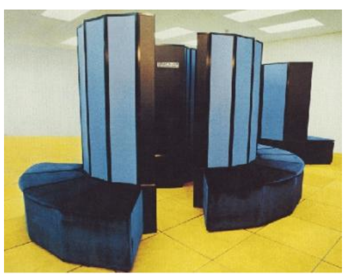

c. 1983
Cray X-MP, SMP
2-4 cores

c. 1988, Cray Y-MP, SGI
Crash

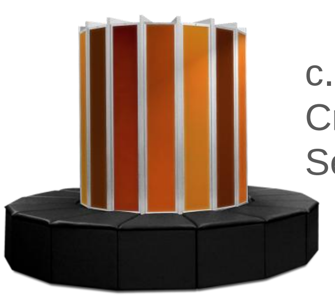

c.1978
Cray-1, Vector processing
Serial

c. 1983, Cray X-MP, Convex
MSC/NASTRAN

Do CAE algorithms scale?

WRF Results on Blue Waters (preliminary)

- WRF V3.3.1
- 1km, 1 bil
- 30 minute
- WSM5 (m
- Results o ... sockets)

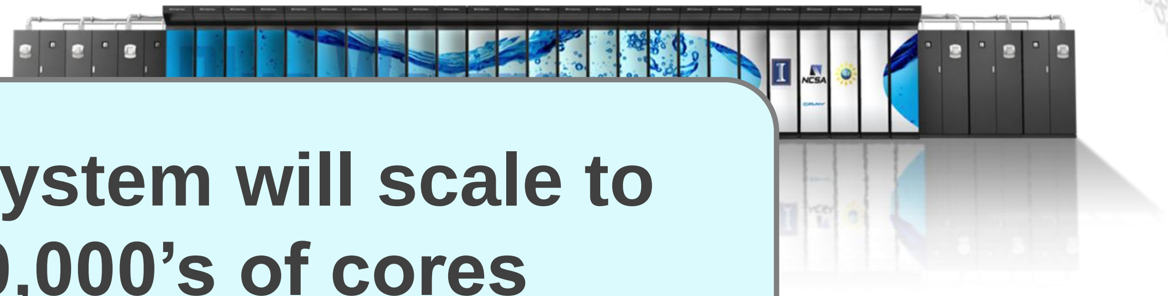

The system will scale to 10,000's of cores

Integer Cores	total nodes	average timestep (seconds)	sustained performance (GFLOPS/sec)	Speedup
2048	64	3.995	2181	1.0
8192	256	1.065	8182	3.8
32768	1024	0.286	30480	15.6
131072	4096	0.142	61485	28.1
262144	8192	0.053	166332	75.4

Cavity Flow Studies using HECToR (Cray XE6)

S. Lawson, et.al. University of Liverpool

- **1.1 Billion grid point model**
- **Scaling to 24,000 cores**
- **Good agreement between experiments and CFD**

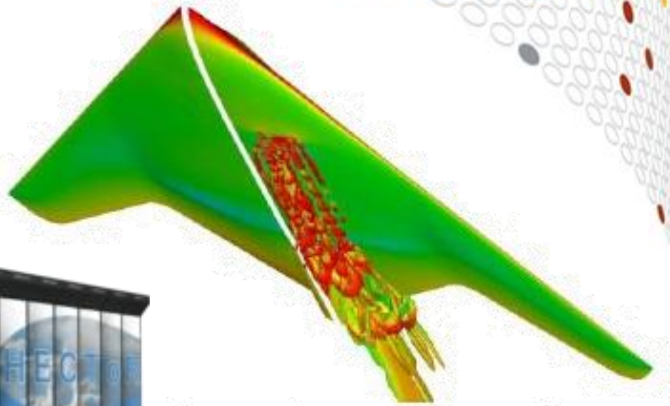

Flow under an Unmanned Combat Air Vehicle (UCAV) with an open weapon bay.

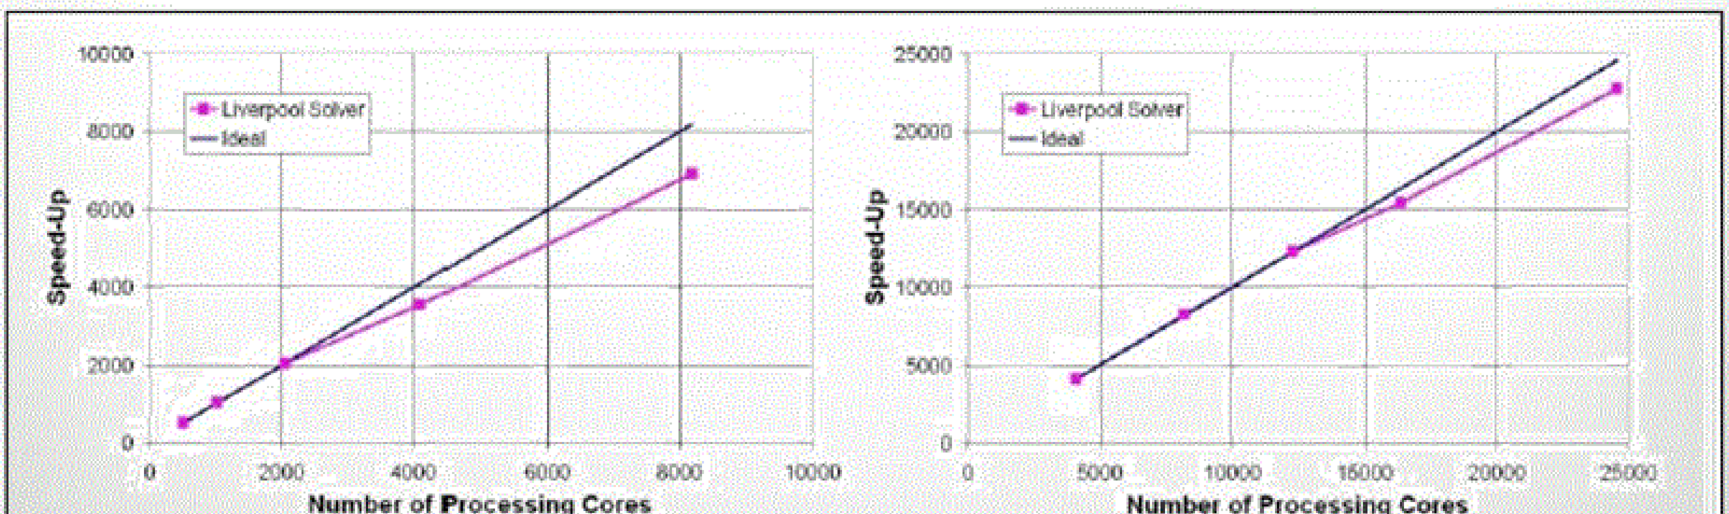

Speed-up of the HMB flow solver using (a) 105 million point mesh with 24576 blocks and (b) 1100 million point mesh with 196608 blocks.

* Ref: http://www.hector.ac.uk/casestudies/ucav.php

CTH Shock Physics

CTH is a multi-material, large deformation, strong shock wave, solid mechanics code and is one of the most heavily used computational structural mechanics codes on DoD HPC platforms.

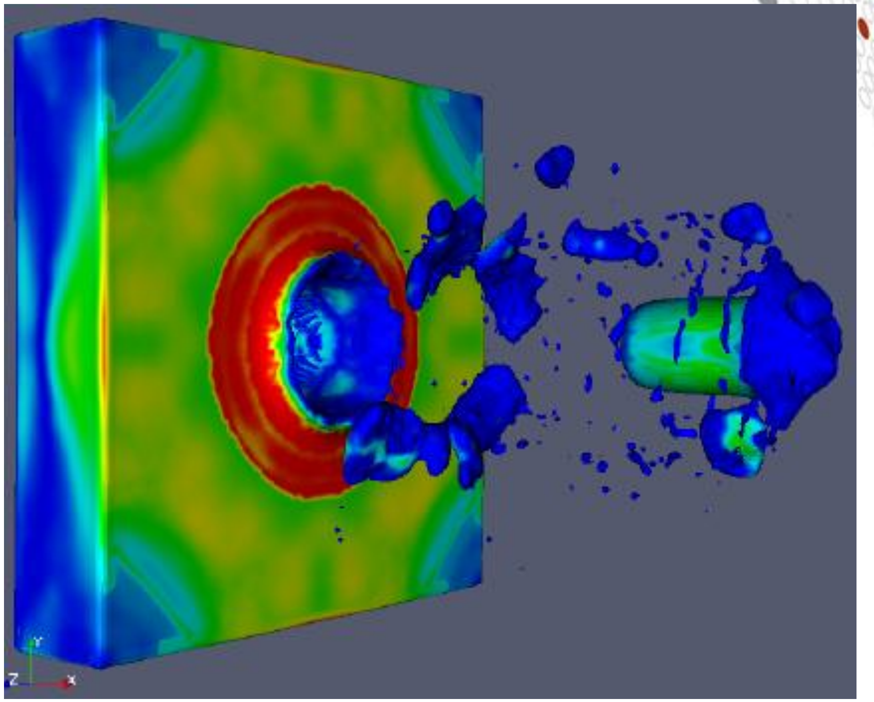

“For large models CTH will show linear scaling to over 10,000 cores. We have not seen a limit to the scalability of the CTH application”

“A single parametric study can easily consume all of the Jaguar resources”

CTH developer

Large Cray systems running ISV applications

- Several of the largest Cray systems are running CAE applications
  - CAE codes scaling to over 10,000 cores
- Both In-house and ISV applications
- Commercial companies are using Cray systems at the HPC centers

Are scalable systems applicable to commercial environments?

Two Cray Commercial Customers

- GE Global Research
  - Became aware of the capability of Cray systems through a grant at ORNL
  - Using Jaguar and their in-house code, modeled the "time-resolved unsteady flows in the moving blades"

    Ref. Digital Manufacturing Report, June 2012
- Major Oil Company
  - Recently installed and accepted a Cray XE6 system
  - System used to scale key in-house code

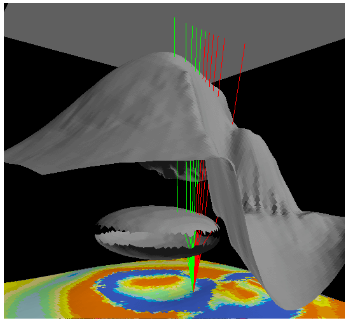

The common thread here is that both of these organizations had important codes that would not scale on their internal clusters

Are ISV applications extremely scalable ?

For many simulation areas…YES!

ANSYS Fluent scaling to >3000 cores on XE6

Aeroelastic Simulation, “Supercritical Wing”

In support of AIAA Aeroelastic Prediction Workshop

Fluent simulation, 13.7 million cells

CRAY XE6, AMD Interlagos 2.1GHz IL16, Cray MPI

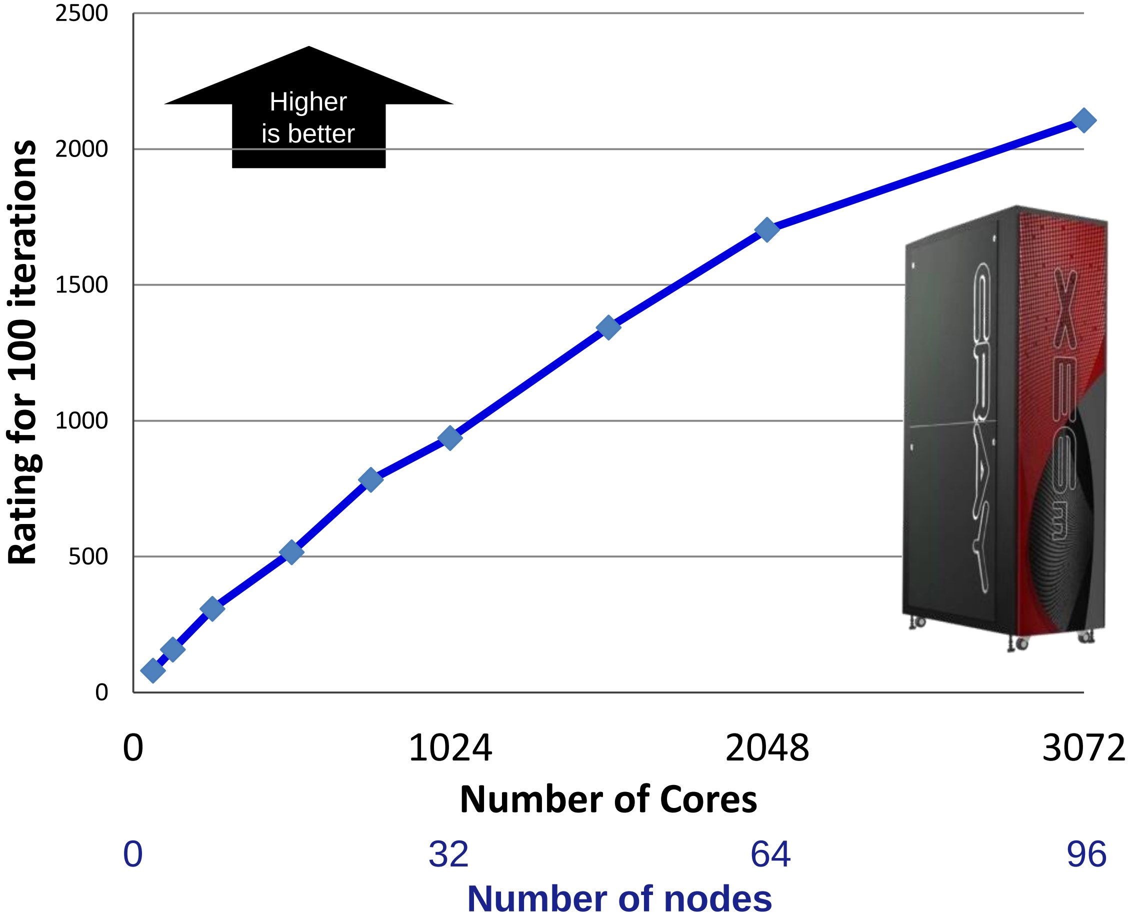

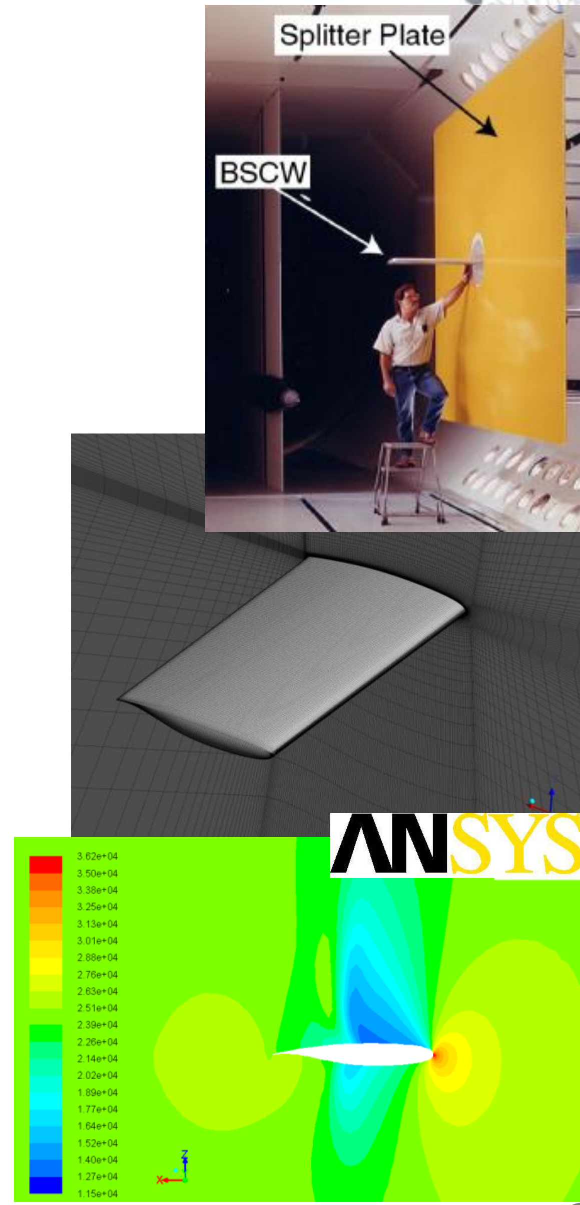

ANSYS Fluent scaling to >4000 cores on Cray XE6

Performance testing of Fluent has shown scalability to over 3000 cores even with this modest size model

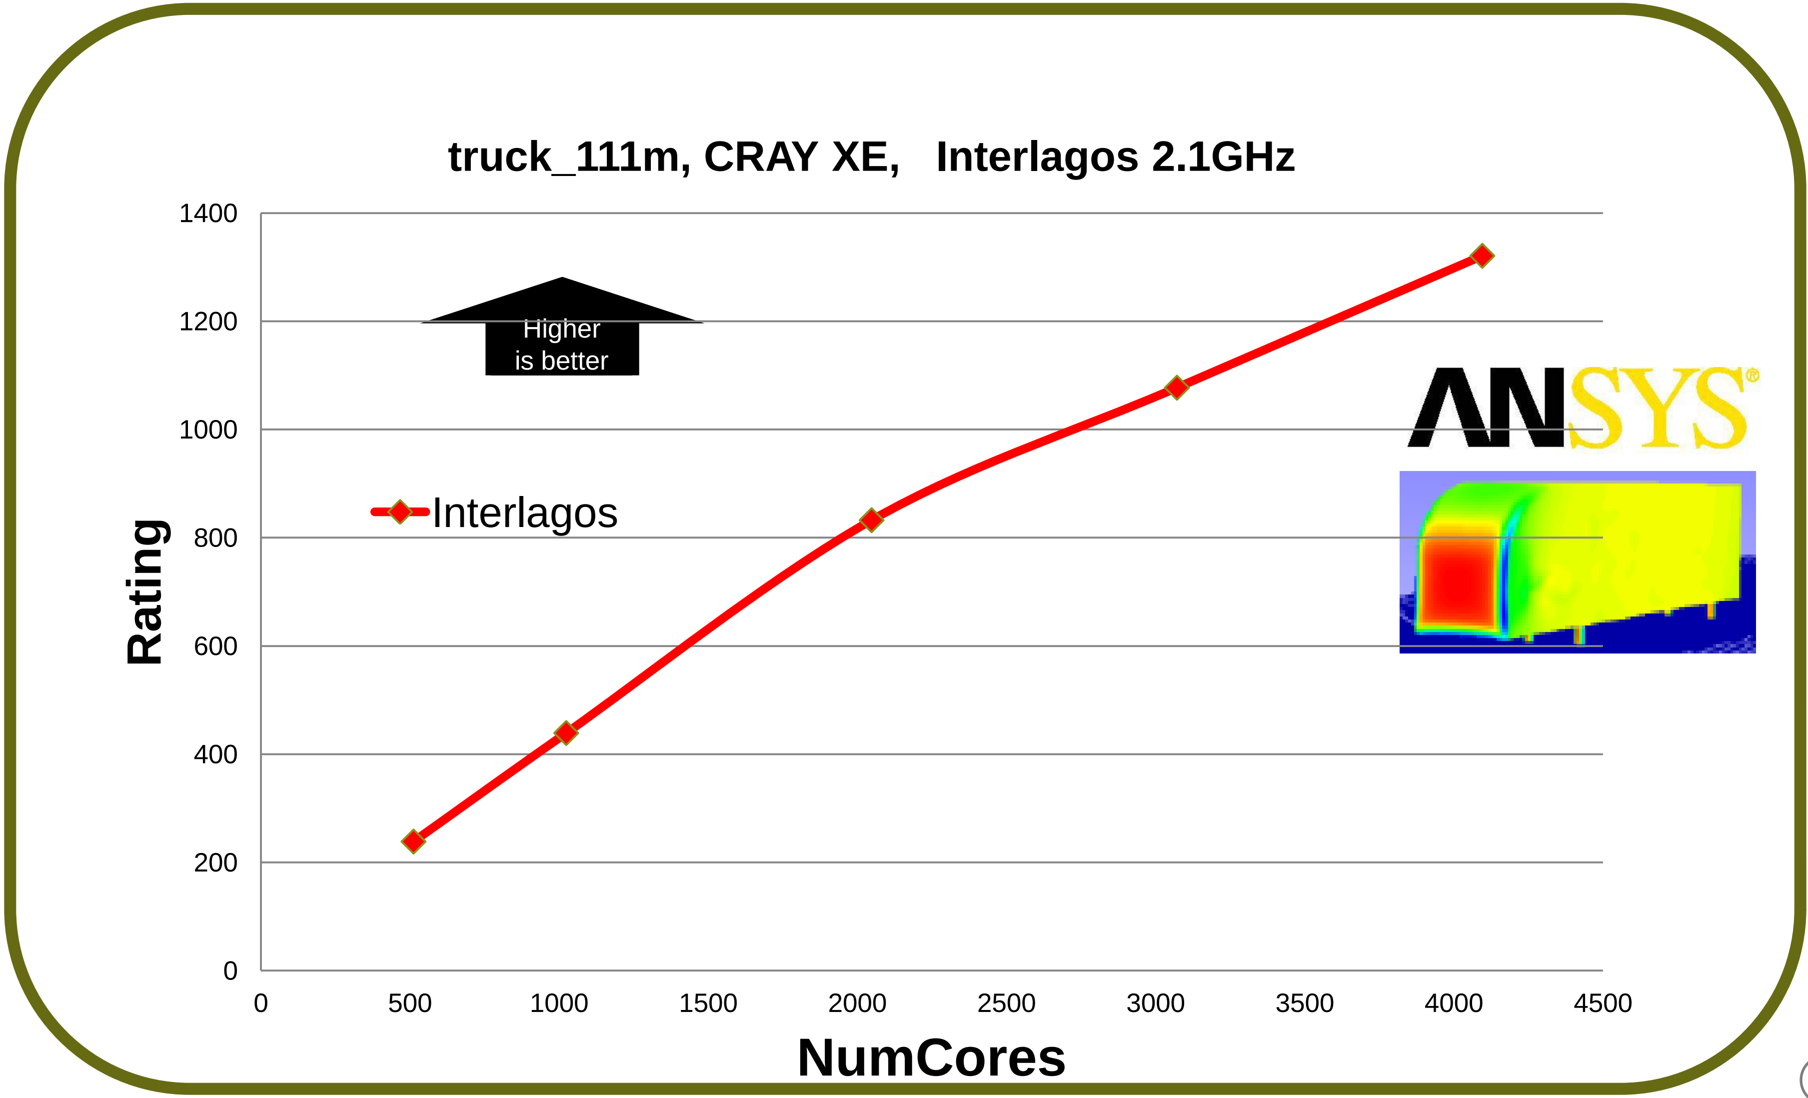

Cray XE6: Extreme Scalability with EXA PowerFLOW

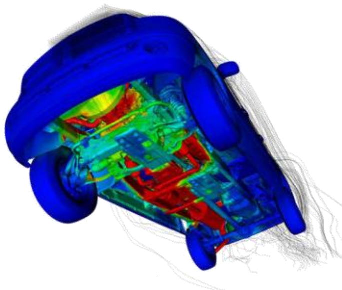

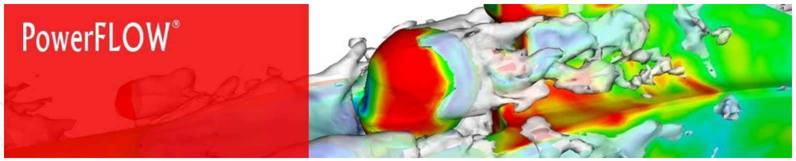

- **Cray XE6: " Scaling to a larger number of cores than any other platform"**
- **Scaling to over 3000 cores**
- **3X the total performance of any other systems**

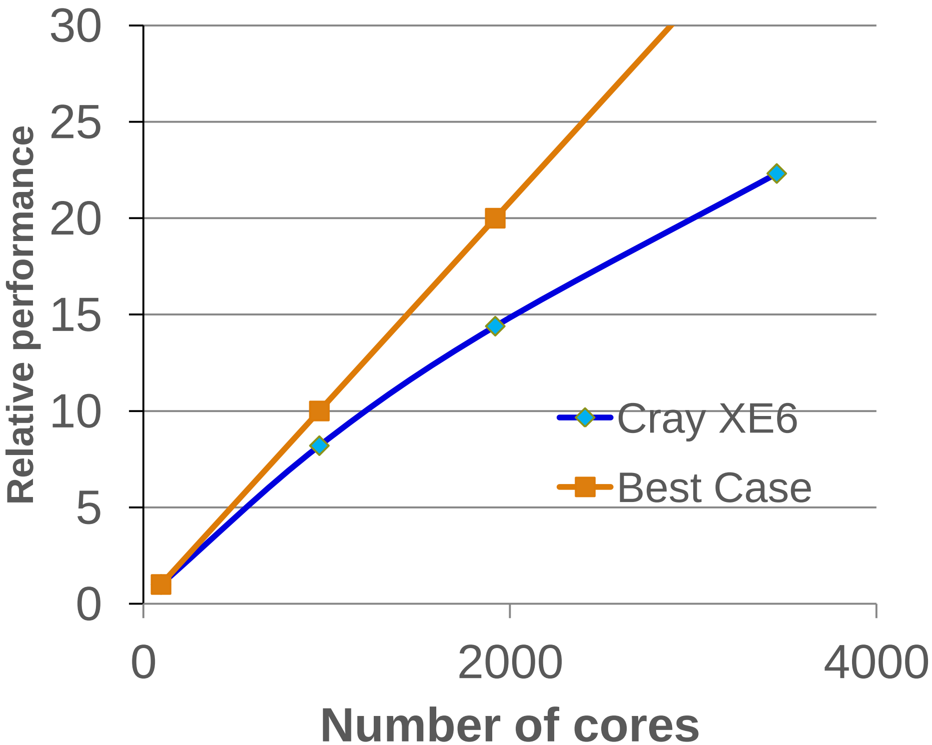

* ref: EXA press release Oct 2011

STAR-CCM+ benchmark
100M cell automotive aerodynamics

Cores/MPI Ranks	Cray XE6, “Interlagos“ 8 “core pairs”, 2.1 GHz Star-CCM+	Speedup
72	24.0 4.5 nodes	1.0
144	12.1 9 nodes	2.0
288	6.0 18 nodes	4.0
576	3.3 36 nodes	7.3
1152	2.0 72 nodes	12.0
2304	1.1 144 nodes	21.2

Performance: Seconds per iteration

Number of Nodes used in that run

LS-DYNA benchmark

Two car crash simulation, 2.4M elements
Hybrid parallel

Total number of cores	Cray XE6 MC-12, 2.1 GHz Hybrid parallel	Speedup
144	21,193 6 nodes	1.0
288	12,274 12 nodes	1.7
576	7,643 24 nodes	2.8
1152	5,258 48 nodes	4.0

Performance: Elapsed time (sec)

Number of Nodes used in that run

Status of select ISV applications

ISV Application	Primary segment	Demonstrated scalability *
ANSYS Fluent	Commercial CFD	>4000 cores
LS-DYNA**	Impact/crash analysis	>4000 cores
CFD++	Aero CFD	>2000 cores
STAR-CCM+	Commercial CFD	>3000 cores
PowerFLOW	External CFD	>4000 cores
RADIOSS	Impact/Crash analysis	>1000 cores
Abaqus	Structural analysis	>64 cores

* Demonstrated scalability typically limited by the simulation model available

** Currently working on a 10M element crash simulation model which should scale much higher

If a model scales to 1000 cores will a similar size model also scale that high?

Not necessarily

Obstacles to extreme scalability using ISV CAE codes

1. **Most CAE environments are configured for capacity computing**
   - Difficult to schedule 1000's of cores
   - Simulation size and complexity driven by available compute
   - This will change as compute environments evolve

2. **Applications must deliver "end-to-end" scalability**
   - "Amdahl's Law" requires vast majority of the code to be parallel
   - This includes all of the features in a general purpose
   - This is an active area of development for CAE ISVs

3. **Application license fees are an issue**
   - Application cost can be 2-5 times the hardware costs
   - ISVs are encouraging scalable computing and are adjusting their licensing models

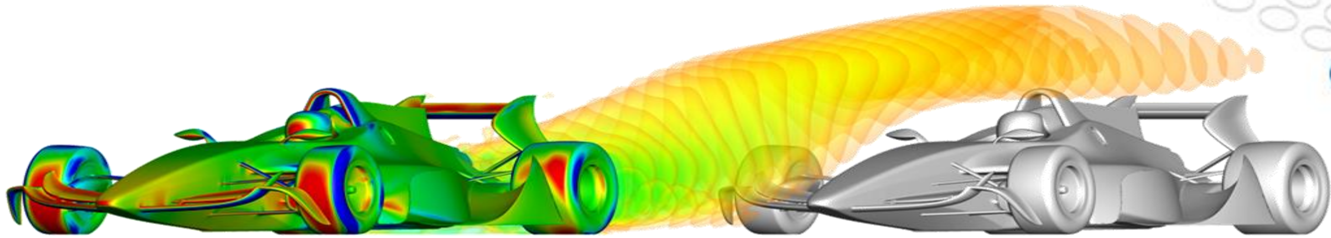

Swift Engineering Inc.

External Aerodynamics

118M cells

unsteady solution, 1350 time steps

moving mesh, rotating boundary condition (tires)

384 cores

350 Hours of elapsed time

Terabytes of data

Cray XE6 with Lustre file system

ANSYS Fluent Scaling on complex industrial model

Pressure based, coupled solver
Compressible, LES

Scaling to 4096 cores with 91% efficiency

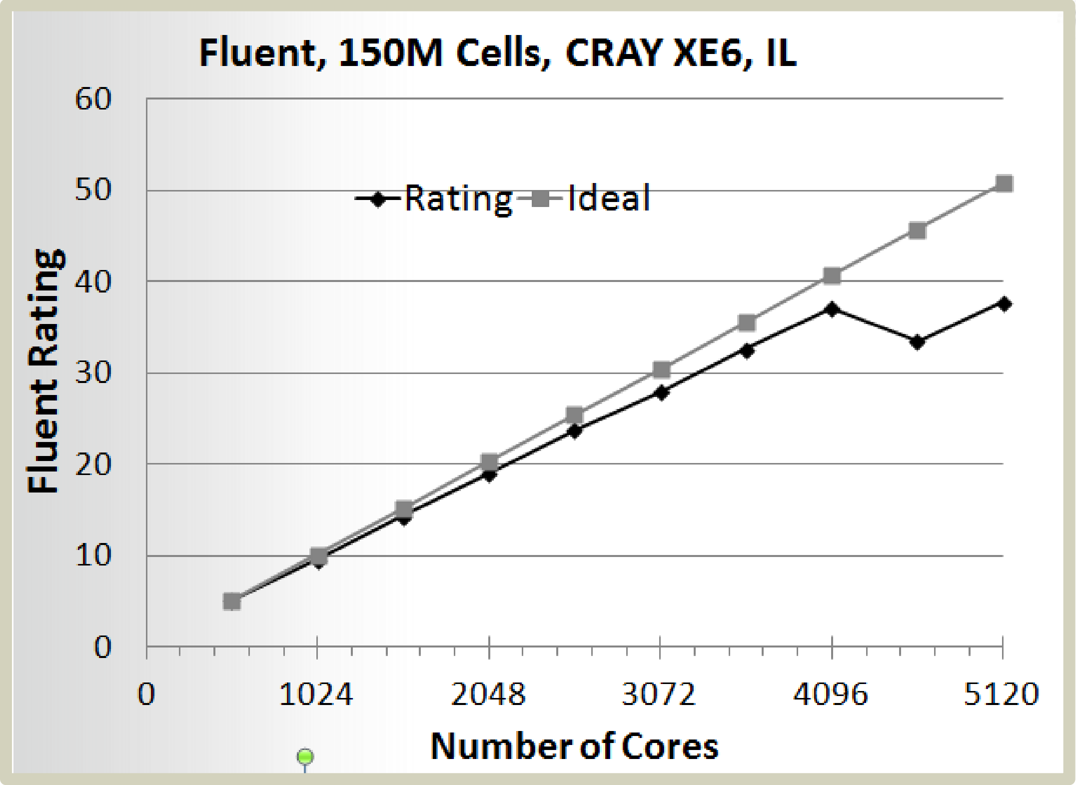

- **Something happens at about 4500 cores but this will be addressed as the project to improve scaling progresses**

- **It is this type of cooperative work between application users, ISVs and Cray, that will lead to extreme scaling for the vast majority of simulations.**

ANSYS presentation at 2012 Cray User Group (CUG)

ANSYS HPC Licensing options

Scalable licensing

- ANSYS HPC (per-process)
- ANSYS HPC Pack
  - Each simulation consumes one or more Packs
  - Parallel enabled increases quickly with added Packs
- ANSYS HPC Workgroup
  - 128 to 2048 parallel shared across any number of simulations on a single server
- ANSYS HPC Enterprise
  - Similar to HPC Workgroup but deploy and use anywhere in the world

Single solution for multiphysics and any level of fidelity

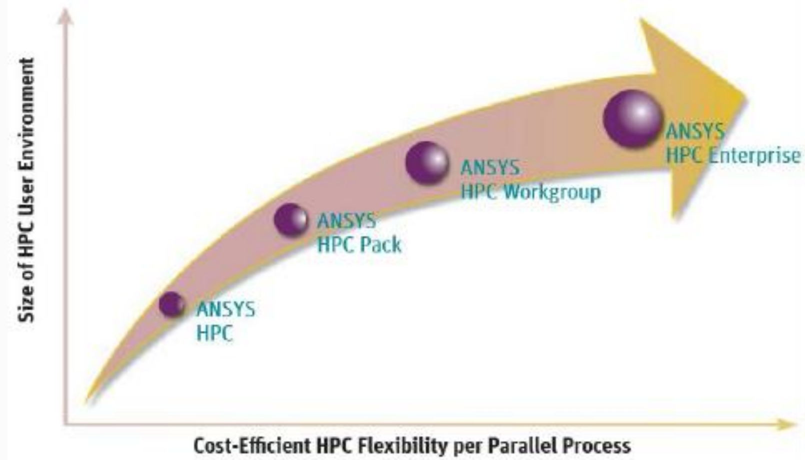

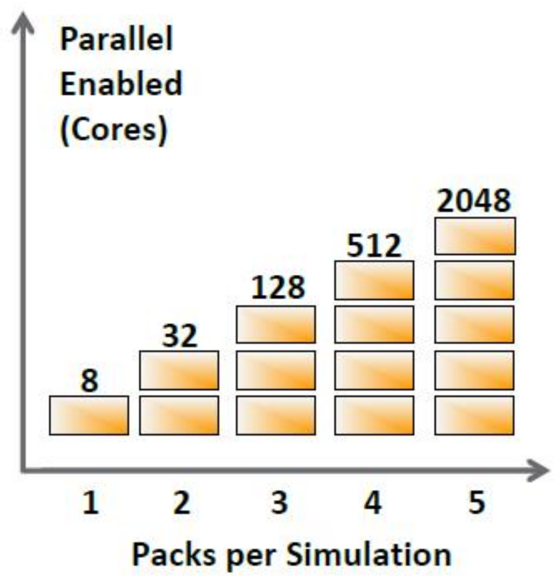

© 2012 ANSYS, Inc.

Summary: Well-balanced Scalable Supercomputer

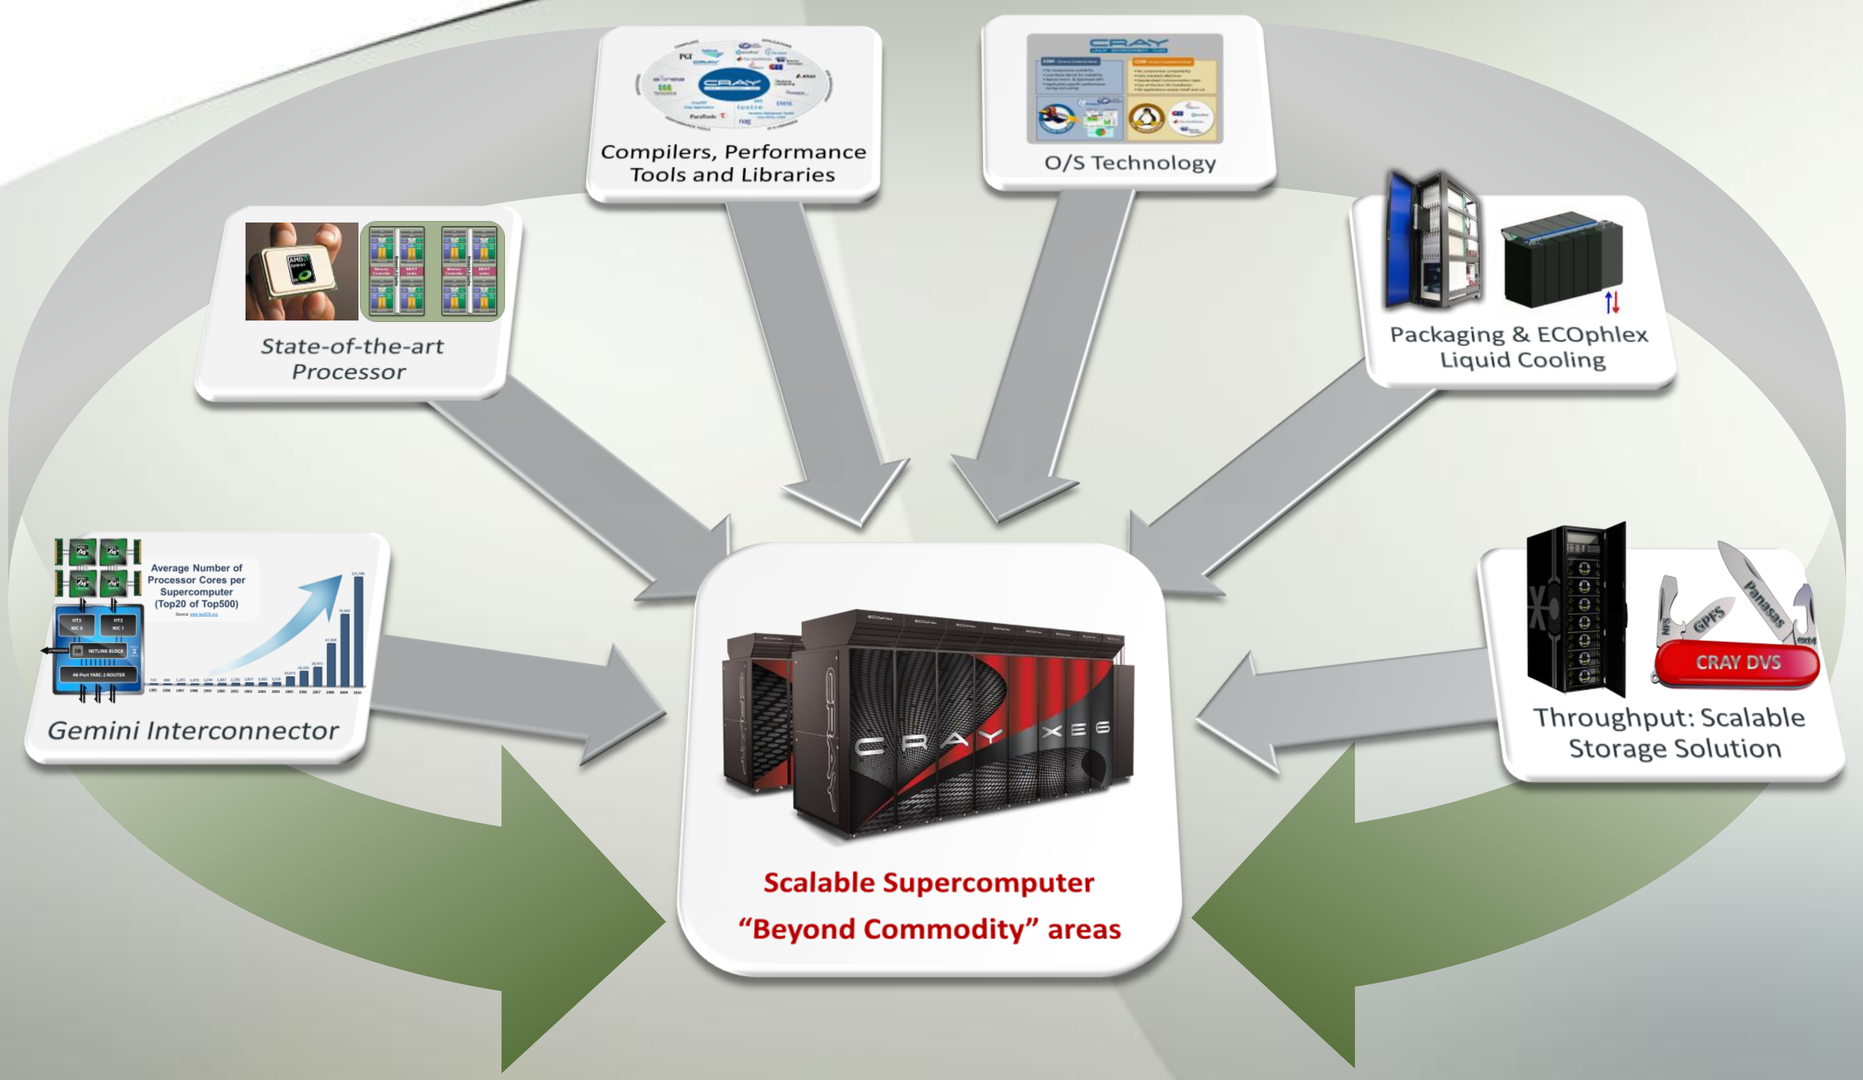

Summary: Well-balanced Scalable Applications Infrastructure

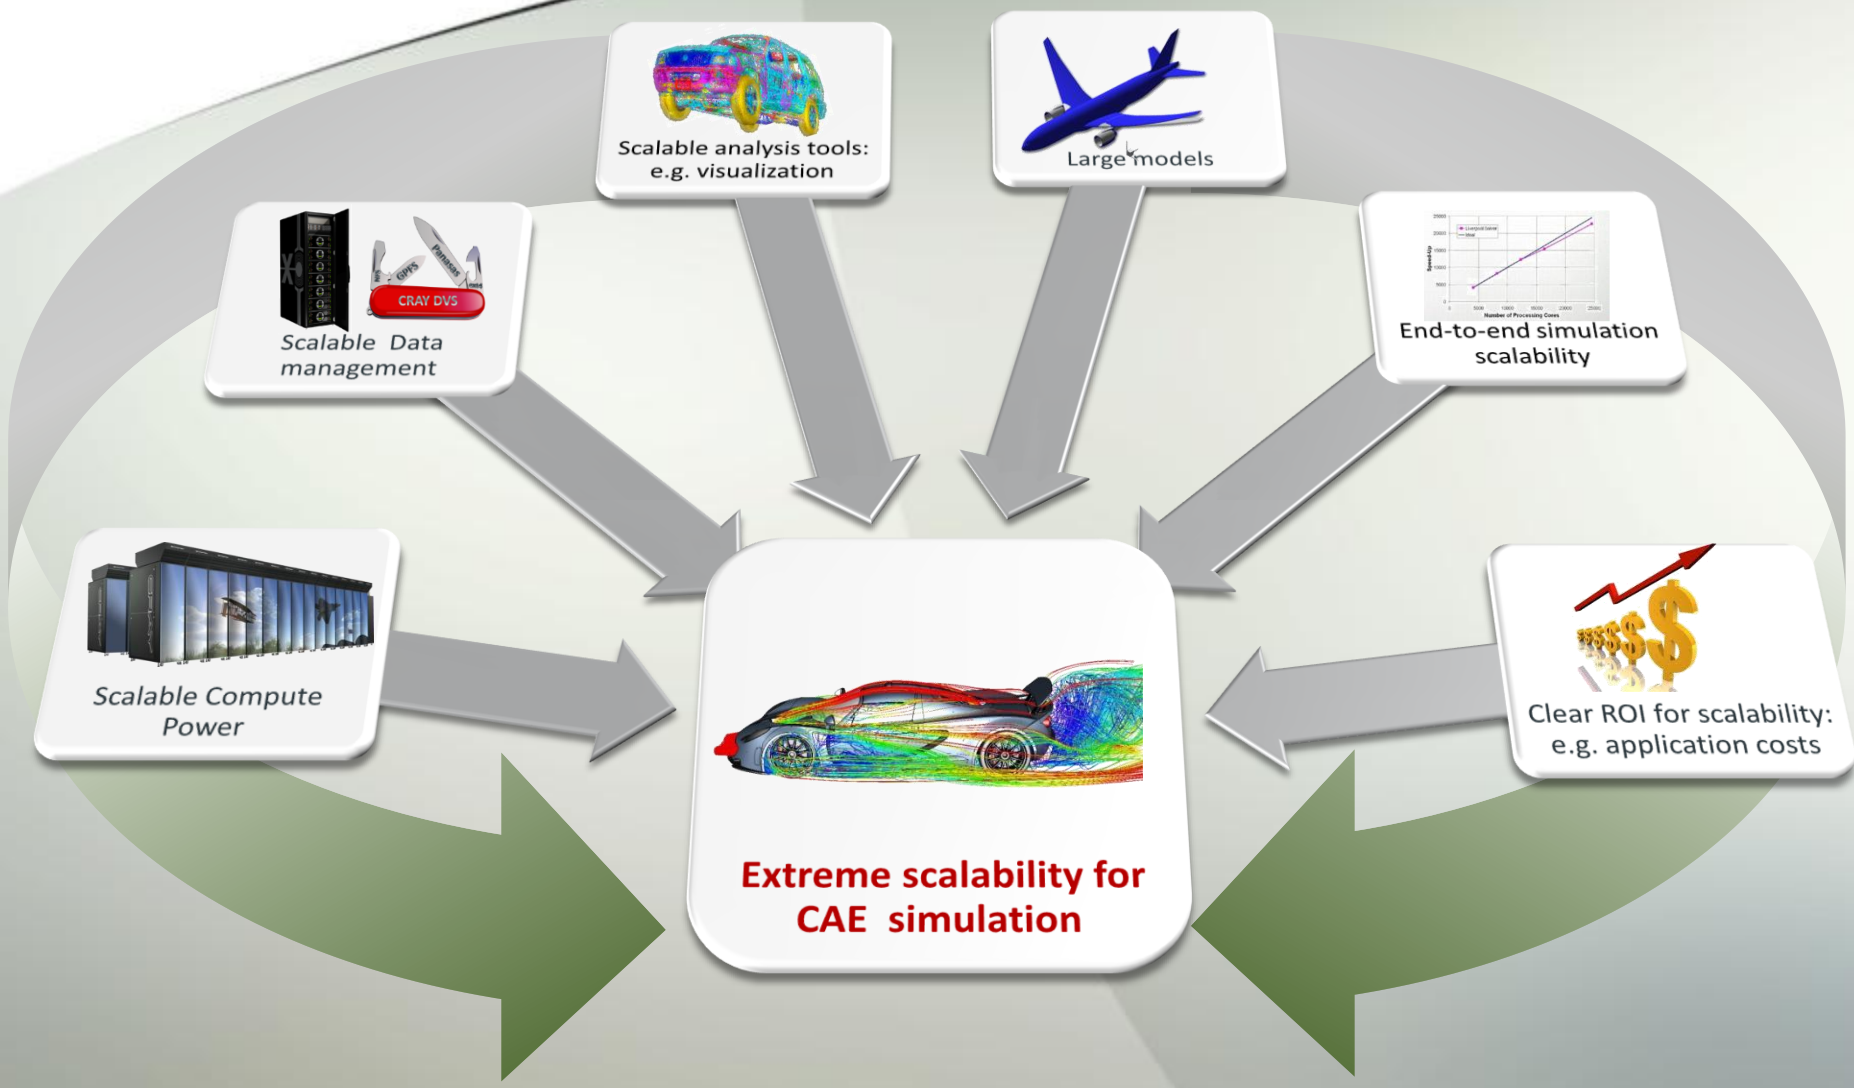

Backup Slides

Random Message MPI Benchmark Test Today Gemini vs. QDR IB

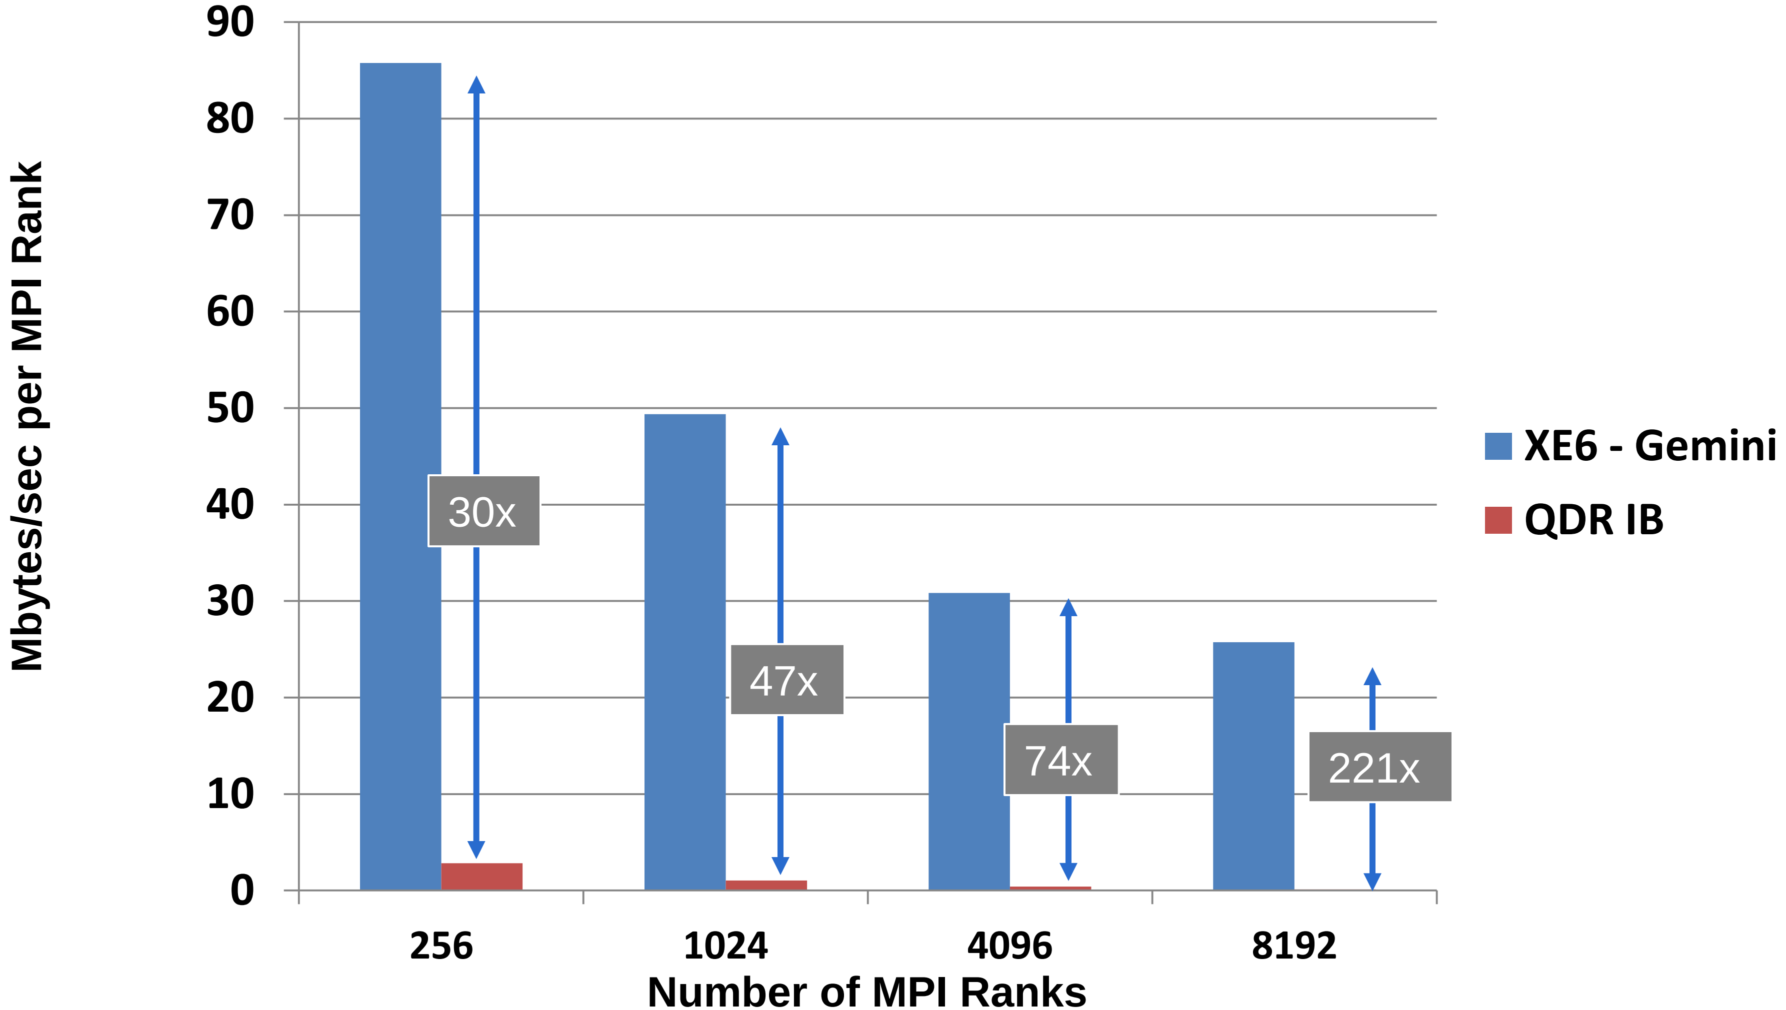

A Comparison of the Performance Characteristics of Capability and Capacity Class HPC Systems

By Douglas Doerfler, Mahesh Rajan, Marcus Epperson, Courtenay Vaughan, Kevin Pedretti, Richard Barrett, Brian Barrett , Sandia National Laboratories